Kingmakers or Cheerleaders?

Party Power and the Causal Effects of Endorsements

Abstract: When parties make endorsements in primary elections, does the favored candidate receive a real boost in her vote share, or do parties simply pick the favorites who are already destined to win? That is, do parties act as kingmakers who swing elections or cheerleaders who merely root on the winning side? To answer this question, we draw on two research designs aimed at isolating the causal effect of Democratic Party endorsements in California's 2012 primary election. First, we conduct a survey experiment in which we randomly assign a party endorsement, holding all other aspects of a candidate's background and policy positions constant. Second, we use a unique dataset to implement a regression-discontinuity analysis of electoral trends by comparing the vote shares captured by candidates who barely won or barely lost the internal party endorsement contest. We find that endorsements do indeed matter, though the magnitude of their impact is smaller than might be estimated from research designs less attuned to recent advances in causal inference.

Thad Kousser
University of California at San Diego

Scott Lucas
San Francisco Magazine

Seth Masket
University of Denver

Eric McGhee
Public Policy Institute of California

Endorsements are a key tool allowing party elites to exert control over the selection of nominees. While American party leaders could once directly pick nominees through secretive caucuses or tightly managed conventions, their power to determine party nominations in the age of the direct primary is negligible unless they have a way to actually influence primary voters. Partisan endorsements, by which partisan elites communicate their preferences directly to donors, political activists, the media, and voters in general, are a central means by which this influence occurs. Parties do play other roles -- by recruiting or discouraging candidates, and by shepherding wealthy donors or energetic volunteers towards their favored candidates -- but a formal endorsement remains the strongest signal that a party can send. These endorsements must exert a real impact on candidates' electoral fortunes if parties are to be kingmakers in American politics.

The influence of prominent endorsements is often taken for granted by political observers, but just how much power do such messages have? If their power over voters is weak, then our political system is essentially a candidate-centered one. The candidate's own skills and appeals would be determinative of both capturing the endorsement and winning the primary, with party leaders essentially relegated to the role of cheerleaders, rooting on strong candidates. On the other hand, if endorsements are highly influential over voters, then party leaders may assume some (if not all) of the authority they once had in the age of the party machines, determining which candidates bear the party mantel and which do not.

In this paper, we take advantage of two unique datasets to offer robust tests of the power of endorsements in party contests. Specifically, we look at the responses by California's Democratic Party to Proposition 14, the 2010 initiative that created a top-two open primary system in that state. The top-two system essentially deprived parties of the power of nominations, placing voters of both parties -- or of no party at all -- in charge of

picking candidates in the June primary. Both major parties, seeking to focus their supporters on a single candidate who would advance to the top-two runoff, responded by creating a robust system of party endorsements. We obtained access to the internal votes by activists that determined the Democratic endorsements, allowing us to see which candidates just barely won or lost the endorsement. We supplement our analysis of endorsements and 2012 election returns with an original survey experiment put in the field during that contest. While the opportunities to study endorsements in this single election are unique, there is no reason to believe that the impact of endorsements on voter behavior in it will be idiosyncratic. Importantly, even though the change in the primary system increased the incentives of strategic party organizations to make endorsements, it did not notably change the reasons why voters might rely on parties to guide them through the process.

We look at the influence of these endorsements in two ways. First, we conduct a survey of 1,000 Californians and present them with a hypothetical ballot for a state Assembly race, randomizing the Democratic Party's endorsement between two Democratic contenders. Second, we perform a regression discontinuity analysis of the results of the 2012 June primary election, using the results of internal party endorsement votes to control for partisan support for the candidates. Both approaches yield similar estimates of the magnitude of a party endorsement -- roughly 10 percentage points of the vote -- although this effect varies considerably based on the attributes of the candidate and of voters.

I. Prior Research Measuring the Impact of Endorsements

Our research focuses specifically on the effect of party endorsements, but follows in the tradition of a larger literature on endorsements by many types of political organizations. This literature is venerable but beset by a number of significant challenges. The main such

challenge is that those who endorse candidates – whether media outlets, celebrities, party elites, or official party organizations – often pick candidates who have a high likelihood of winning in the first place. This makes the endorsement endogenous to the candidate's electoral appeal. Partially, this is a result of endorsers simply being strategic; there is little value in backing a loser, whereas a winner who perceives a benefit from the endorsement can express her gratitude in tangible ways once in office. This is also due to the fact that the same qualities that make a candidate appealing to endorsers also make her appealing to voters and donors. An endorsed candidate may do better than her rivals, but it is difficult to know whether that vote share difference was caused by the endorsement or whether both were caused by factors related to candidate quality.

Early studies of the power of newspaper endorsements in elections tended to ignore measures of candidate quality or campaign spending, but still found important trends. McClenghan (1973) examined newspaper endorsements in local elections in Texas, finding that those with endorsements did better than those without, a result echoed in Erikson's (1976) study of county-level returns in 1964. Gregg (1965) noted in a study of endorsements in California elections that endorsements seemed more effective in local elections than in statewide ones. Hooper's (1969) study of state legislative elections in Illinois found that endorsements did not seem to convince voters to cross party lines, but could explain a great deal of variation within a party's vote:

> [W]hen we ask what is the major source of variation that distinguishes one Democrat from another and one Republican from another the answer is newspaper endorsements. (Hooper: 305)

Why might endorsements actually move votes? Quite simply, they provide information that voters lack, helping them learn about candidates and make a complex

decision with relatively modest effort (Kahneman, Slovic, and Tversky 1982; Lupia 1994; Nisbett and Ross 1980). Such cues are imperfect, of course – novice voters can be led astray from their preferences by endorsements (Lau and Redlawsk 2001) – but they are still highly useful. This information is particularly helpful in environments in which party labels are not available, such as primary elections or many local contests (Stein and Fleischmann 1987). In general elections, conversely, endorsements may have little effect, particularly in periods or on issues on which the electorate is already deeply divided (Gimpel 1998). That said, Ladd and Lenz (2009) find substantial conversion effects from newspaper endorsements in the 1997 British parliamentary elections, suggesting that endorsements can matter even in high-profile contests where the partisan consequences are clear and the party structure has been stable for several generations.

Arceneaux and Kolodny (2009) examine interest group endorsements in the 2006 Pennsylvania statehouse elections, finding that low-information Republican voters use endorsements by liberal interest groups as a negative voting cue, providing them with key information about the candidates that they otherwise lacked. This echoes a finding by McDermott (2006) that AFL-CIO endorsements in congressional elections can help a candidate among Democratic voters but hurt her chances among the Republican electorate.

Several studies focus specifically on primary elections, under the assumption that endorsements – particularly by party elites – are more critical to such contests' outcomes. These studies are of central importance to the understanding of political parties. Party elites, after all, cannot directly pick candidates to represent their parties. Since the advent of the direct primary more than a century ago, party leaders have sought to choose their nominees by communicating with loyal party voters in primary contests. They do this largely through

4

the use of endorsements, but also through the strategic allocation of funding and campaign expertise (Bawn et al. 2012; Cohen et al. 2008; Masket 2009).

Studies of elite control over parties tend to assume that endorsements have some effect on primary voters, and several have found such effects. Dominguez (2011), for example, provides a detailed analysis of congressional primaries, finding that partisan endorsements can increase a candidate's vote share, even controlling for funding and candidate quality. She determines that a five percent increase in the share of partisan endorsements is associated with a one-point increase in the primary vote share. Steger's (2007) study of endorsements in presidential nomination contests finds that a candidate's share of party endorsements made prior to the Iowa caucuses positively predicts her vote share in subsequent presidential primaries, even controlling for the candidate's spending, position in the polls, and quality. And Stone, Rapoport, and Abramowitz (1992) estimate that the AFL-CIO's backing of Walter Mondale was worth an additional six percentage points of the vote in the 1984 Iowa Democratic caucus.

Cohen *et al.* (2008) have compiled perhaps the most thorough dataset of partisan endorsements in presidential nomination contests. Their analysis echoes Steger's; a candidate's share of key endorsements prior to Iowa does a better job predicting primary performance than polling, fundraising, or media coverage. In their estimation, a ten percent increase in the candidate's share of partisan endorsements (properly weighted) corresponds to an eight-point increase in that candidate's share of delegates won in state primaries (p. 292). The authors are forthright, however, in pointing out the considerable potential for selection bias (p. 257-9), since a lack of early endorsements might have kept these candidates out of the race when they would have performed perfectly well otherwise.

The central concern to estimating the impact of an endorsement, however, is endogeneity. Masket (2009), for example, found that elite endorsements correlated with higher vote shares for candidates in local primary elections in California after controlling for candidate quality and spending. However, candidates with electoral appeal and good early fundraising totals were better positioned to capture these endorsements, making it unclear whether the endorsers were acting as kingmakers or simply rallying behind the strongest candidate. It is difficult to control for candidate quality, because no comprehensive measure of everything that confers electoral strength is available. Most researchers tend to follow Jacobson and Kernell's (1981) approach of treating those who have previously held elective office as high quality and those without such experience as low quality. This approach can be effective, but it is unlikely that any single variable alone can capture all the ways in which candidate quality varies, many of them ineffable and elusive to quantitative measurement.

II. California's Adaptation to the Top-Two Primary

Our research design uses two approaches to break the endogenous relationship between a candidate's underlying electoral strength and whether she wins the endorsement. Both leverage unique data sources from California's new top-two primary.

California's legislature has developed a reputation as one of the most polarized legislatures in the country (Masket 2009), and systematic comparison to other states has shown that this reputation is well-deserved (Shor and McCarty 2011). As a result, many political observers, reformers, and officeholders have converged on the goal of mitigating party polarization in the state's politics. Toward this end, reformers have pushed two specific reforms over the past few years – redistricting reform (placing a bipartisan panel of

citizens, rather than the state legislature, in charge of redrawing legislative districts) and the top-two primary. It is this second reform that motivates our analysis.

The top-two primary -- employed only by Louisiana, Washington, and now California -- is perhaps the most "open" form of primary rules. Not only can all voters participate, but they may choose among all candidates for office, not just those within one party. Furthermore, one possible outcome of a top-two primary is two candidates of the same party advancing to a runoff election. Theoretically, the more moderate of the two candidates would be well positioned to win such a contest, as she could win the support not only of some of her fellow partisans, but also of independents and perhaps most of the voters of the other party.

This reform was widely seen as threatening to the establishment of both of the state's major political parties. After all, it essentially placed party nominations in the hands of voters not registered with the party, and as the U.S. Supreme Court remarked in 2000, "a single election in which the party nominee is selected by nonparty members could be enough to destroy the party" (*California Democratic Party v. Jones* 2000). The language of Proposition 14 specifically redefined the June contest as a "voter-nomination primary election" rather than a party contest, reflecting the fact that this election was no longer one in which a party could determine its nominee.

The major parties moved quickly to respond to this shift in the laws by capitalizing on their party endorsement system, enabling them to reassert some control over party nominations and converge voter support around candidates more faithful to the party banners. Endorsements would be considered and voted on at party county and state central committee meetings. The state Republican party actually moved first on this front (Van Oot 2011), but the Democratic party soon followed suit. Under the new Democratic rules, any

candidate may be considered endorsed with a 70 percent vote at a county central committee meeting, or at the state central committee meeting with a 50 percent vote for incumbents or a 60 percent vote for challengers.[1] The list of endorsees would then appear within the official ballot pamphlet mailed by the Secretary of State to voters prior to the election. Notably, indications of endorsements did not appear on the sample ballot pages, but rather in a later page of the pamphlet where they were much less visible. (No endorsements appeared on the ballot itself, making the information available to voters in this election similar to what is generally available to voters in contests in other years and in other states.) Candidates winning the endorsement would also be free to trumpet it far and wide in their campaign communications.

To be sure, the 2012 primary was an unusual one for candidates, party elites, and unusually attentive voters, who were navigating a form of election that was rare for the nation and entirely new for the state. But for the lay voter, it was mostly business as usual. The primary was held in June, as it typically is in California, and was conducted with much the same fanfare. Voter turnout was on the low end of normal. Candidate information, including party endorsements, were available in state election literature for those voters who sought it -- just as similar information was in previous primary elections -- and voters still had to choose among a range of candidates for each office.

Just how much influence did these endorsements have? In our theory section, we lay out the potential paths through which they might affect candidate vote shares. In the following sections, we describe the results of two efforts to determine the power of the party endorsements. The first method is through the use of a survey in which we proposed two fictitious Democratic candidates for an Assembly seat and randomized the party's endorsement. The second method is through a regression discontinuity analysis of the actual

election results of the June 2012 primary election. Each of these approaches to causal inference breaks the endogeneity that plagues a simple comparison of the performance of candidates winning and losing endorsement battles. In our survey experiment, we randomly assign the endorsement, holding the candidates' other attributes constant. Our regression discontinuity compares the vote shares of candidates who just barely won and narrowly lost an endorsement, but who are otherwise quite similar in their electoral strength. We focus our analysis in both sections on competing Democratic candidates, because Democrats are the dominant party in California[2] and the one for which we have obtained internal records of endorsement votes.[3]

III. Theory: Why Should Endorsements Matter?

In this section we briefly ground our empirical expectations in prior work on political behavior, explore additional implications of our main hypothesis, and clarify what the evidence from each our two tests can show about how party endorsements may help candidates.

First, we note that our survey experiment and our regression-discontinuity design will teach slightly different lessons. The survey will show whether endorsements can help candidates by doing nothing more than sending a signal to voters. We experimentally manipulate this signal and look for its impact, holding all else equal. By contrast, our regression-discontinuity design, drawn from observational evidence about whether an endorsement given in early 2012 benefitted candidates in June of that year, gauges the total effect of endorsements through their signal to voters, donors, and other potential endorsers. Endorsed candidates may win higher vote shares because party's imprimatur helps to

coordinate potential donors in an area or because it leads to a crescendo of other endorsements from interest groups and local political leaders.

An endorsement can benefit candidates in June through these causal paths even if voters pay no attention to the party's signal itself. If our regression discontinuity shows that candidates who just barely win the party endorsement see a jump in their vote share, this could be because of voter signaling effects alone, broader campaign effects alone, or some combination of the two. By contrast, our survey experiment isolates the impact of the signal that an endorsement sends to voters. In the remainder of this section, we specify exactly how we expect that signal to influence voters.

First, we expect that when a candidate receives the party endorsement in our survey, the candidate should receive a higher level of support from respondents than when that candidate is not given the endorsement. This is a straightforward hypothesis that works directly through voter information, though of course the information that an endorsement provides for voters could lead to higher vote shares through multiple avenues. It could give voters a cue that experts who share their interests know that this candidate will best represent them. As experiments conducted by Lupia and McCubbins (1998) and Boudreau (2009) have shown, those are conditions under which voters are likely to follow cues. It could tell them that they will be following other voters. Learning about an endorsement might inform voters that a candidate will be the focal point for their co-partisans, leading them to follow along with the group's coordination. Finally, it could tell them that they will be voting with elites. McClosky, Zaller, and Chong (1985) show that receiving and comprehending the elite viewpoint may push voters to internalize the norm as their own, with voters in the 2012 primary acting in much the same way as the voters who adopted elite positions on an anti-gay California proposition in the 1980s.

Regardless of how the signal might work – by pushing voters to follow a cue, to follow other voters, or to follow elites – an endorsement should have an independent causal impact on a respondent's likelihood of supporting a candidate. If we observe this effect in our randomized experiment, it will be evidence that parties indeed have the power to be kingmakers.

> *Hypothesis 1. Endorsements should benefit candidates. In our survey experiment, a respondent will be more likely to support a candidate who (by random assignment) has won the endorsement, while in our observational data, candidates who barely win the endorsement should earn a higher primary vote share than those who barely lose it.*

Second, we ask which voters should respond to the signal of an endorsement. For strong Democrats, an endorsement by the Democratic Party should have a strong effect, while for Republican voters it should help not at all (or perhaps even hurt). This follows the logic that Zaller (1992) lays out: when political elites disagree, as they consistently do in partisan elections, partisans follow their own party's cues. The divergence in the impact of an endorsement should grow with a voter's adherence to her party affiliation.

> *Hypothesis 2. The impact of a Democratic Party endorsement should be strongest for a strong Democratic adherent, weaker for leaning Democrats and independents, and non-existent or negative for Republican registrants.*

Third, we ask which types of candidates will benefit most from winning the Democratic Party's endorsements. Although Democratic lawmakers in California are generally liberal in their voting positions (Masket 2007; Shor and McCarty 2011), there are internal divisions between traditional Democrats who are often associated with labor interests and less liberal Democrats who are referred to as "business Democrats," "BizDems," or "The Mod Squad." We attempt to distinguish between these two types of candidates in our observational study by recording whether a Democratic candidate listed a business profession on her official ballot occupational designation, and in our survey build this difference into our two candidate biographies. We do so to test whether or not an

11

endorsement means more for one group than the other. Perhaps traditional Democrats benefit the most from an endorsement because it can fit into the voter's existing "schema," her frame of reference based on prior expectations about that candidate which makes her more readily acceptant of information that confirms this stereotype (Fiske 1986). On the other hand, it could be that an endorsement only has an impact when it causes a voter to update her prior beliefs. If this is true, learning that a "Business Democrat" has won the party endorsement despite holding a less liberal position could be real news to the voter, and have a bigger impact on her behavior.

> *Hypothesis 3a. The Democratic Party endorsement should have its largest effect on candidates with a* **traditional** *background.*

> *Hypothesis 3b. The Democratic Party endorsement should have its largest effect on candidates with a* **business** *background.*

IV. Evidence from a Survey Experiment

Survey experiments allow researchers to randomly assign some respondents to a control group and others to one or more treatment groups, holding all else equal and thus isolating the causal impact of the treatment (Sniderman 1996). Our treatment here is the party endorsement, which we can give to a particular candidate for a subset of our sample. While work by Iyengar et al. (2001) applied this technique to test the effects of endorsements on proposition voting and Brader and Tucker (2009) and Brader, Tucker, and Duell (2012) probed the impact of party endorsements of policy positions, we have not seen any prior research conducting a survey experiment on the effect of party endorsements on candidate vote shares.

We conducted our experiment by contracting with the polling firm YouGov, which administers online polls and uses weighting to ensure that the sample reflects general

population characteristics. Our poll went into the field from May 29 through June 9[th], 2012, just around the time of the June 5[th] statewide primary, and included a sample of 1,000 California registered voters. All but 23 reported that they had voted or planned to vote in this election. At the third question in a poll that went on to ask respondents a series of policy questions, we outlined a hypothetical state Assembly race with three candidates. The field included a "traditional Democrat" whom we named Greg Johnson, a "business Democrat" whom we named Sam Guthrie, and a Republican we named David Robertson. We randomized the question so that 1/3 of respondents were told that Johnson had won the Democratic Party's endorsement in this race, 1/3 were told that Guthrie captured the endorsement, and the remaining 1/3 were not given any endorsement signal at all. After presenting the biographies, we asked respondents for whom they would vote. The set of options that we provided to respondents at the end of the question was designed to mimic the actual ballot format as closely as possible. A sample of an actual ballot from this election can be seen in the Appendix, Figure A1.

The full text of the question follows:

"As you may know, California will be holding its primary election on June 5[th]. Your state Assembly district has its own set of candidates, but if you were choosing between the three candidates described below to represent you in the legislature, how do you think you might vote?

Greg Johnson, a Democrat, is an experienced educator who began his career as a classroom teacher and now serves on his local school board. He favors increased funding for schools in California, as well as tough environmental regulations and an increase in the minimum wage. [1/3 OF SAMPLE: **Johnson won the official endorsement of the California Democratic Party in this race**.]

Sam Guthrie, a Democrat, is an environmental entrepreneur who founded a solar energy installation company. He wants to build California's green economy by broadening consumer incentives, and to keep our communities safer through strict gun control laws and a community-based approach to policing. [1/3 OF SAMPLE: **Guthrie won the official endorsement of the California Democratic Party in this race**.]

David Robertson, a Republican, owns and operates a small business that provides many jobs in his community. He wants to cut taxes and red tape in Sacramento to reduce unemployment across the state, and will use his experience as a volunteer with the sheriff's department to pursue tough-on-crime policies.

On the ballot, each candidate will be listed, along with his party preference and occupation. If this was your ballot, which candidate would you support?

> [] GREG JOHNSON
>   My party preference is the Democratic Party
>   Educator
>
> [] SAM GUTHRIE
>   My party preference is the Democratic Party
>   Business Owner
>
> [] DAVID ROBERTSON
>   My party preference is the Republican Party
>   Business Owner"

Figure 1 shows how respondents distributed their votes among the candidates, for each of our three randomized conditions. (All of the vote shares reported in this paper incorporate survey weights.) When neither Democrat was given the endorsement, Greg Johnson captured 29% of the overall vote, while Sam Guthrie's support was 19%. Yet when a different group of respondents was told that Greg Johnson was the Democratic Party's endorsee, he did markedly better, winning 36% of the vote. This seven percentage point difference was statistically significant at the 95% confidence level, a clear sign that an endorsement can help a Democrat with a traditional career background like Johnson's. For Guthrie, though, winning the endorsement was less beneficial. For the final 1/3 of voters who were told that Guthrie was the endorsee, his vote share increased by only one percentage point, to 20%, and the difference was not statistically significant.

Figure 1

The main results of this survey experiment provide partial support for Hypothesis 1. Both candidates performed better when they had won the endorsement, but the magnitude

14

of the effect was much larger for the traditional Democrat than it was for the Democrat with

a business background. This provides support for Hypothesis 3a, suggesting that

endorsements mean more for some types of candidates than other.

Figures 2 and 3 explore the contention of Hypothesis 2 that the impact of an

endorsement will be stronger for voters who are more strongly associated with the cue-giver,

the Democratic Party. Figure 2 shows the vote shares for each candidate in each condition,

but only looks at voters registered with the Democratic Party. It shows effects that are

stronger. Greg Johnson captured 59% of the Democratic vote when no endorsement was

made, but 69% when he was the endorsee. Sam Guthrie won only 37% of Democratic

support in the control group – evidence that he was not seen as the traditional candidate at

the heart of his party – but climbed to 41% support when he won the endorsement.

Figure 2

Figure 3 divides Democrats more finely. It draws on a 7-point measure of party

affiliation, allowing us to separate weaker partisans from stronger ones. The data in the

graph comes from a logit regression that predicts the impact of Greg Johnson winning the

endorsement on the likelihood that a respondent will vote for him, holding constant a set of

demographic controls[4] in case they were not perfectly balanced by our randomization of the

endorsement treatment. The two lines on the graph report the predicted likelihood of

support when there is no endorsement and when Johnson is endorsed, with the gap between

the lines signifying the endorsement effect. The average estimated effect of this

endorsement across all voters in the model was ten percentage points (an effect that is

significant at the 95% confidence level). It is clear, though, that this effect grows stronger

the more closely a voter is affiliated with the Democratic Party. There is almost no effect

(three percentage points) for Republicans, while a real gap emerges for true independents in

15

the middle of the scale (nine points) and grows further – indicating a stronger endorsement effect – as the graph moves to independents who lean Democratic, to weak Democrats, and finally to strong Democrats (for whom the endorsement effect is 15 percentage points). This provides clear support for Hypothesis 2.

Figure 3

The bottom line of our survey experiment is that the signal that an endorsement provides to voters about a candidate can exert a strong and significant causal impact, but only under certain conditions. It benefits a traditional Democrat in our study, but the boost it gives to a Democrat with a business background falls short of significance. For both types of candidates, the impact of the endorsement is stronger among voters who are more closely allied with the Democratic Party.

V. Evidence from a Regression-Discontinuity Design

Our second approach to the question of endorsement effects is to use "regression discontinuity" (RD) techniques to analyze election outcomes from the 2012 election in California. The regression discontinuity (RD) design has seen a tremendous burst of popularity in recent years, long after its initial introduction by Thistlethwaite and Campbell (1960). The essence of the design is fairly straightforward. Many decisions to implement a treatment are based entirely on the value of a continuous "forcing" variable.[5] Cases falling above a predetermined threshold receive the treatment and those falling below it do not, while the goal of the analysis is to evaluate the effect of the treatment on some outcome of interest. The classic example is an academic merit award (the treatment) which is handed out based on a student's performance on a qualifying test (the forcing variable), with the goal

of identifying the effect of this award on the student's later academic outcomes (Thistlethwaite and Campbell 1960).

The challenge is that the forcing variable often reflects a key concept—in the example above it is scholastic aptitude—that has its own effects on the outcome variable. In that case, any difference between the treatment and control groups might simply reflect the impact of that concept and not of the treatment per se. The RD design circumvents this problem by comparing only those cases falling just above and just below the threshold on the forcing variable. Those who score just above and just below the threshold are otherwise quite similar: much of the observed effect of the award can be attributed to the treatment itself rather than to the impact of the forcing variable. If the chosen bandwidth around the threshold is sufficiently narrow, then the act of crossing the threshold becomes effectively random. Any discontinuity in the relationship between the forcing and outcome variables at the threshold can therefore be attributed to the effect of the treatment alone (Lee and Lemieux 2009).

In our application, we compare the share of the vote won in the June 2012 primaries by candidates who, about six months before that election, just barely won or fell a few votes short of capturing the Democratic Party's endorsements. Candidates clustered around this threshold, we show, are – aside from some winning the endorsement and some not – quite similar, allowing us to isolate the impact of the endorsement. We thus draw all of our causal inferences from these cases rather than from the fortunes of the hopeless candidates who failed to win any support in the internal party competition for an endorsement or from the incredibly strong candidates who won 100% of the internal party vote. To estimate the causal effect of the endorsement, we use two approaches: the regression-based procedure laid out by Imbens and Kalyanaraman (2009), and the randomization inference approach

17

first developed by Fisher (1935) and applied to RD designs by Cattaneo et al. (2013). Both approaches yield the same substantive findings.

Critical to this analysis, then, is the internal party process used to determine endorsements. As noted earlier, the party used a multi-stage process for deciding endorsements, first at "pre-endorsement conferences" around the state where local party activists could vote on the endorsement, and then at a statewide convention if no candidate received enough votes in the first stage. The precise threshold for awarding an endorsement differed depending the stage of the process and whether the candidate was an incumbent, but formal rules governed this variation. We use this continuous vote to examine the impact of endorsements on vote share in the June election using the RD framework.

RD's key identifying assumption of randomness near the threshold is violated if certain individuals are able to manipulate the forcing variable to ensure a score just above the threshold. For example, if skilled and experienced candidates were able to game the endorsement process to ensure a favorable outcome—perhaps by stacking the voting committee with supporters, for instance—then those barely receiving an endorsement would be better candidates than those falling just short, and the treatment might reflect nothing more than this process of self-selection. Manipulation of this kind does not create problems for RD so long as the candidates cannot *precisely* control the forcing variable (Lee and Lemieux 2009; McCrary 2008). As long as there is a sizable element of the process that is beyond a candidate's control, outcomes near the threshold can be treated as effectively random and the identifying assumption is preserved. Fortunately, it is possible to test for random local assignment by examining whether cases above and below the threshold are similar in their baseline covariates, and by testing for a higher density of cases just above the threshold. Either might indicate the presence of manipulation and cast doubt on the causal

estimate. This is a problem we take very seriously, since it is easy to imagine candidates attempting to manipulate the outcome and succeeding at it. We will test our results for these sorts of limitations below.

To further probe our data, we employ a randomization inference approach to our RD analysis, as pioneered by Cattaneo et al. (2013). This approach treats the observations near the threshold as a local randomized experiment and is particularly optimal for small sample sizes. It runs 10,000 simulations of the test statistic (a difference of means, in this case) using the data near the threshold, and reports the range of results as a confidence interval. In this way, it focuses particular attention on the internal validity of the causal effect (Keele, McConnaughy, and White 2012). We report the results of both the traditional RD and the randomization inference RD below.

As mentioned above, our forcing variable is each candidate's share of the party's vote to endorse in a given race. We center this vote around the threshold for the endorsement in the final deciding vote, so the threshold always falls at zero. Our outcome variable is the Democratic candidate's share of the vote cast for all Democratic candidates in each district. The top-two primary raised the prospect that voters would cross party lines in large numbers to support candidates on the other side of the aisle, but the incidence of such cross-party voting was probably small: the share of the primary vote in each district that went to Democrats and Republicans was broadly predictable from elections when almost no crossover voting was allowed (McGhee and Krimm 2012), and at least one survey estimate suggested only 17 percent of voters voted outside their party identification (Ahler, Citrin, and Lenz 2013). Thus, most voters were likely choosing among candidates of the same party and the relevant comparisons concern how much of the Democratic Party vote each

Democratic candidate received. We run our results with different outcome variables to be certain of our results.

The results of this basic RD design can be found in Figure 4 and in Table 1. The bandwidth around which we estimate the regression discontinuity is determined by Imbens and Kalyanaraman's (2009) mean squared error method and is relatively wide—plus or minus 12.3 percentage points around the win/loss threshold. In this simple first look, there is a clear discontinuity at the cut point.[6] The discontinuity is 15.1 percentage points, very similar to the size of the effect we found in the survey data among Democratic voters. This effect is consistent with our survey findings, moves in the direction that we hypothesize, and is significant at the 92% confidence level in a one-tailed test (or 83% in a two-tailed test). The estimated impact of an endorsement is fairly sensitive to the size of the bandwidth – the discontinuity is 6.2 percent for a bandwidth half the size, and 10.6 percent for a bandwidth twice the size – but is consistently positive and large enough to have a substantive impact on election outcomes. For instance, former Democratic Assemblymember Lori Saldana fell just short of winning her party's endorsement for the 52nd Congressional District, capturing 59.5% of the vote at the state convention rather than the 60% that she needed. In the primary, she also fell agonizingly close, finishing third with 22.1% of the vote while Democrat Scott Peters advanced to the general (where he unseated incumbent Brian Bilbray) with 22.6%. An endorsement boost as large as any of those we estimate at different bandwidths would have made a clear impact in this and other tight races.

Figure 4

Table 1

Is the apparent effect of an endorsement the product of selection bias? In particular, are strong candidates gaming the endorsement process to ensure their own success? We

address this question in two ways. First, we look to see if there is a discontinuity in the density of the forcing variable itself (Caughey and Sekhon 2011; McCrary 2008).[7] If there are significantly more candidates falling just above the threshold than just below, it suggests that some candidates may be able to manipulate the process when the outcome is likely to be close. The results of this validation can be found in Figure A2 in the Appendix. There is no clear discontinuity in the distribution at the threshold: the log difference in the height of the density at that point is 0.35, with a standard error almost three times as large (0.91). As one would expect based on Figure 4, there is a high density of cases at a winning margin of 30 percent, which is the margin for almost every candidate who received 100 percent of the endorsement vote.[8] But the distribution looks almost uniform otherwise.

The second test involves examining baseline covariates within the same RD framework. The logic is similar: if the outcome is manipulated, then candidates above and below the threshold should differ in terms of variables important for success in the primary. If this is the case, then those variables could provide alternative explanations (aside from the endorsement) of why endorsed candidates capture a greater vote share. Even if there were no manipulation, with our small sample size it is especially important to determine whether any discontinuities are a product of chance alone. Considered this way, it is akin to validating that a randomization procedure was properly conducted in an experimental trial (Lee and Lemieux 2009).

We can identify two important covariates that predate the endorsement vote and might account for the differences. The first is incumbency (Caughey and Sekhon 2011): are there more incumbents just above the threshold than just below it? The answer is a qualified yes, though as with all of these analyses the sample size is small so the effects are difficult to estimate precisely. Within the bandwidth estimated in Figure 4, one of six candidates was an

incumbent below the threshold, compared to three of eleven above it. Thus, it is possible

that at least some of the effect we have found is due to incumbency alone. That said, if we

limit that analysis to non-incumbents, as in Figure 5, the size of the discontinuity actually

grows slightly to 19.6 points (p=0.18) and loses some of the apparent shift in slope that

characterized the results with incumbents included.

Figure 5

Another significant pre-treatment covariate is the amount of money a candidate has

raised prior to the endorsement vote. Strong fundraising can signal to those voting on the

endorsement that a candidate has broad support among activists and interest groups and will

be able to raise the resources necessary for a credible campaign. It can also serve as a proxy

for the sort of general skill set required to be an effective candidate. The key is to measure

this money *before* the vote so it is not contaminated with the effect of the endorsement itself.

Figure 6 shows this discontinuity estimate. We have deviated the fundraising totals

from the mean for each chamber (state Assembly, state Senate, and U.S. House) to reflect

the different amounts typically raised for each type of race. The result suggests that, if

anything, those candidates above the threshold raised slightly *less* money than those below it,

though the difference is small. This result stands in contrast to the result with votes, where

the endorsed candidates perform somewhat better. Interestingly, there is also no clear

difference in the money raised *after* the endorsement vote (see the figure's second panel).

The amount is perhaps slightly higher and the variance does seem to increase somewhat, but

there is no clear and consistent effect. This helps to clarify the causal path through which

endorsements work. We see no evidence here that endorsements coordinate donors on the

most viable candidate or help a candidate raise more than she otherwise would.

Endorsement effects, then, likely work through signals to voters (the sort of impact measured in our survey experiment) or through campaign dynamics other than fundraising.

Figure 6

An alternative way of estimating the effect of an endorsement, which is appropriate for studies like ours with relatively small numbers of observations, follows the Cattaneo *et al.* (2013) randomization inference RD method. For this approach, we selected two different window sizes for the endorsement vote variable: 15 and 20 points above or below the endorsement margin. These windows are admittedly arbitrary, but they are close to that selected by the Imbens method (see above), and it is plausible to consider any vote among a small number of party leaders close if a candidate is within 20 percentage points of the victory margin.

The top line of Table 2 reports the results of this method for each different window size. The independent variable here is the margin of support received by the candidate in the endorsement convention, and the dependent variable is the candidate's share of the Democratic vote in the June primary. The outcomes reported are the estimated effect size, its p-value, and its confidence interval.

Table 2

As can be seen in the first line, the randomization inference approach yields an estimated effect of 11.7 percent of the vote for the ±15-point window, with a p-value of 0.07. Again, this estimated effect size is very close to that calculated by our other approach to estimating the RD effect and in the survey experiment. The effect for the ±20-point window, meanwhile, is a slightly-larger 13.1 points, and that is statistically significant at the p≤.05 level. These findings provide additional confidence that the estimated effect is a real one and not simply due to chance. We can further test for selection bias within the

randomization inference framework just as we did using the Imbens- Kalyanaraman approach. The lower rows in Table 2 show the results for district partisanship and pre- and post-endorsement fundraising. None remotely approaches statistical significance, giving us increased confidence that the effect we have calculated on primary results is real.

The core results presented here are robust to a number of alternative measurement and analysis decisions. We coded all Democratic candidates according to their occupation as listed on the ballot, and found, consistent with our survey results, that the effect of an endorsement appeared to be stronger for traditional Democrats than for those listing a business occupation on the ballot (see Table 1). While this is in some ways counterintuitive – one might have guessed that the endorsement would matter more if it gave voters new information about "business Democrats" who do not fit the traditional party mold – it is robust across our two types of data. It provides support for the idea we set forth in our theory section that voters may be most receptive to endorsements when these signals confirm the schemas and expectations that voters already hold.

In another robustness check, we find that the results are very similar when limited only to uncompetitive seats. We also excluded candidates who were left uncontested when their opposition at the endorsement stage ended up dropping out by the election stage. The discontinuity was smaller without these cases (12.7%, p=0.25), suggesting that at least part of the endorsement's power is to scare off opponents. Nonetheless, the effect was still positive.

VI. Conclusion

As much prior research has shown, endorsements at first glance appear to confer a major electoral benefit upon candidates. In California's 2012 primary, candidates who won the Democratic Party's endorsement captured an average of 86% of the party's vote,

24

compared with an average of 32% for candidates who did not win the endorsement. Even in a regression that controls for incumbency status and a district's party registration, the estimated effect of an endorsement is 40 percentage points. But this apparent effect alone is not enough to tell us whether parties act as kingmakers or mere cheerleaders in primary battles. The candidates who win endorsement battles at party conferences and conventions at the beginning of a campaign may be those who were already destined to perform well on Election Day. Researchers need to approach the question of causal inference carefully to determine whether parties simply endorse only the strongest candidates or whether their endorsement adds to the electoral strength of those who earn it.

This paper takes two approaches to this thorny empirical question. Through a survey experiment and a regression discontinuity design, we find qualified evidence that an endorsement does indeed exert a causal impact on candidate vote shares. The effect is not nearly as strong as one sees at first glance, averaging an estimated 10 to 15 percentage points. The impact is strongest among voters who closely associate with the Democratic Party, in keeping with past findings about how individuals follow party elites (Zaller 1992). In our survey experiment, a randomized endorsement had a statistically significant impact on the vote share of a traditional Democrat, but though it helped a Democrat with a business background, its impact was smaller and fell short of significance. In our regression-discontinuity, the estimated endorsement effect had a similar magnitude and approached or surpassed significance at the 95% confidence level depending on the estimation technique used.

What do our results suggest for broader questions about the nature of party power? Given that the evidence that we present here comes from just one state, just one election, and just one party, we should be cautious about generalizing to other contests and venues.

Our approach here, rather, is to sacrifice some measure of external validity in our attempt to maximize internal validity through close attention to research design. Yet there is little reason to believe that this case is particularly idiosyncratic in the impact of endorsements on voters. As we have argued, the top-two primary gave parties a strong incentive to make endorsements, but did not significantly alter the voting experience of lay voters. They could seek out endorsement information if they wanted, as in any previous election, but otherwise it was a relatively typical election for them.[9] If anything, our use of the California case biases our results against finding an impact of party endorsements, since the parties had even less control over the choices offered to voters than usual. Further research will be necessary to see if the effect that we estimate holds across all legislatures over time, but we lay out an approach that can help guide wider exploration.

What is truly unique about this election is the data we have obtained: internal party records of endorsement decisions and a survey experiment, both for measuring the same phenomena. Using these sources, we see that endorsements matter, but that when their impact is scrutinized carefully it becomes clear that they are not determinative in all races. The 40-50 percentage point performance gap between candidates who win and who lose endorsements is mostly a function of their underlying strength, while our experiment and RD design suggest that 10-15 points of this gap may come from the true independent effect of winning the endorsement itself. Parties do act as cheerleaders for their strongest candidates, while ignoring those who seem destined to lose. Yet when party elites wade into a closely contested endorsement fight in a district, it matters whom they chose in that internal battle. Putting the party's mark of imprimatur on one of two otherwise strong candidates, our survey and observation evidence shows, can make a notable difference. In this constrained way, parties can use their endorsements to be kingmakers.

26

Table 1. Traditional Regression Discontinuity Results

	Estimated Treatment effect	p-value	Confidence interval
All Democrats	15.12	0.07	[-6.61, 36.84]
Traditional Democrats	11.18	0.32	[-10.97, 33.33]
Business Democrats	0.77	0.99	[-112.47, 114.00]

Notes: Estimated treatment effects are the impact of the endorsement on the share of the primary vote for Democratic candidates, using the method outlined in Imbens and Kalyanaraman's (2009). "Business Democrats" are candidates who listed a business occupation on the ballot, and "Traditional Democrats" listed other occupations.

Table 2: Randomization Inference Regression Discontinuity Results

Dependent variable	Window of ±15			Window of ±20		
	Treatment effect	p-value	Confidence interval	Treatment effect	p-value	Confidence interval
Share of Democratic primary vote	11.72	0.07	[-1.25, 24.78]	13.12	0.03	[0.74, 25.63]
Democratic registration in district	0.02	0.63	[-0.06, 0.12]	0.02	0.66	[-0.05, 0.1]
Percent of Democratic fundraising pre-endorsement	0.16	0.31	[-0.17, 0.48]	0.11	0.43	[-0.16, 0.38]
Percent of total fundraising pre-endorsement	0.05	0.81	[-0.36, 0.46]	-0.03	0.87	[-0.33, 0.28]

Figure 1. Survey Experiment Results For All Respondents
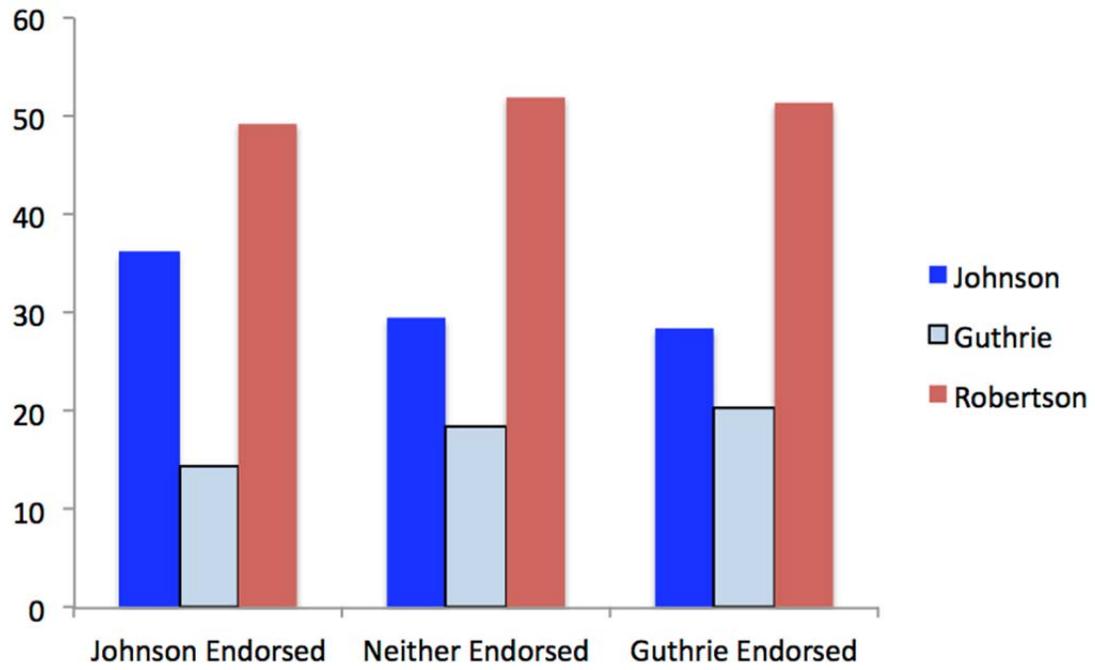
Figure 2. Survey Experiment Results, Democratic Respondents Only
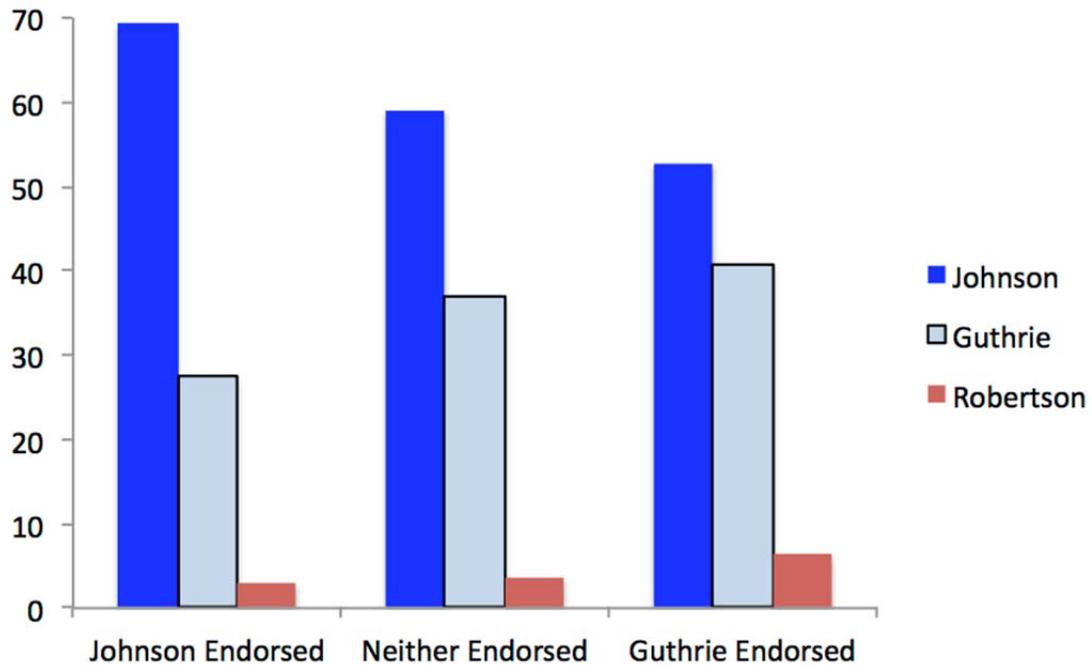
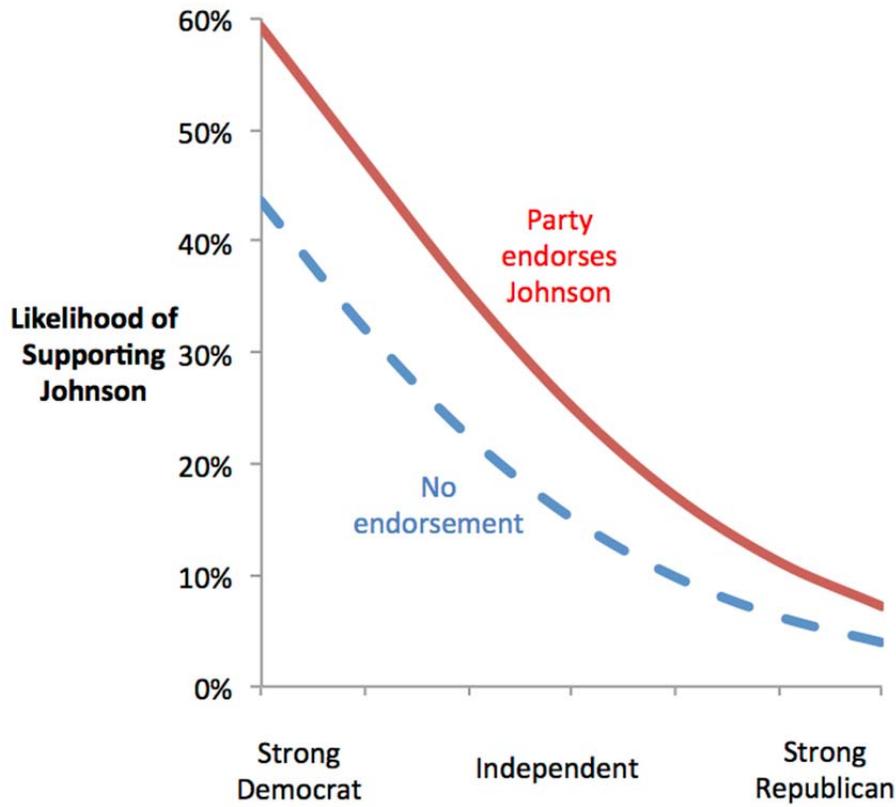**Figure 3. The Interaction of Endorsement with Party Identification**

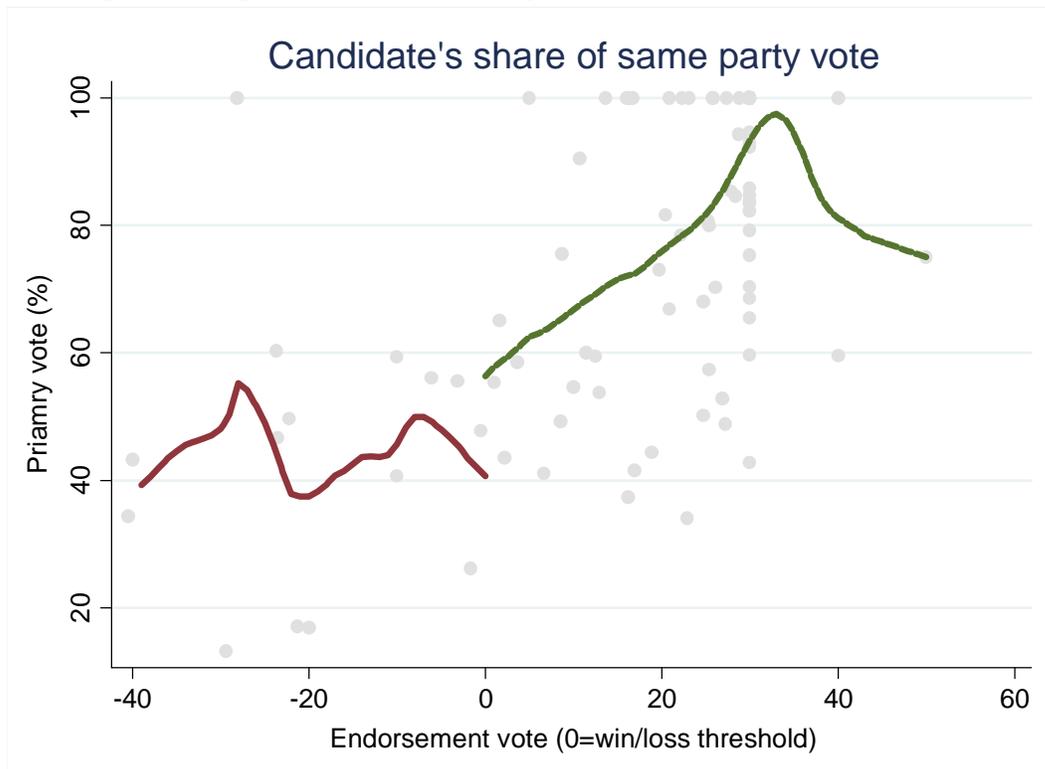**Figure 4. Regression Discontinuity: Effect of Endorsement on Vote Share**

Candidate's share of same party vote

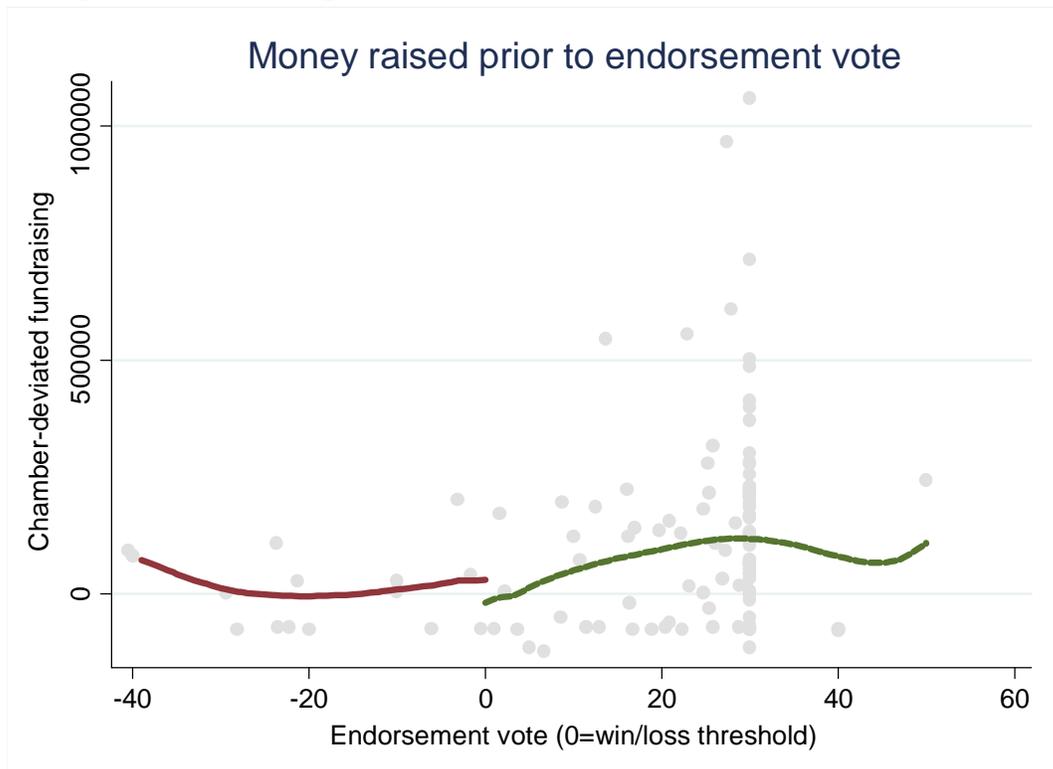**Figure 5. Regression Discontinuity: Effect of Endorsement on Non-Incumbents**

Candidate's share of same party vote, non-incumbents only
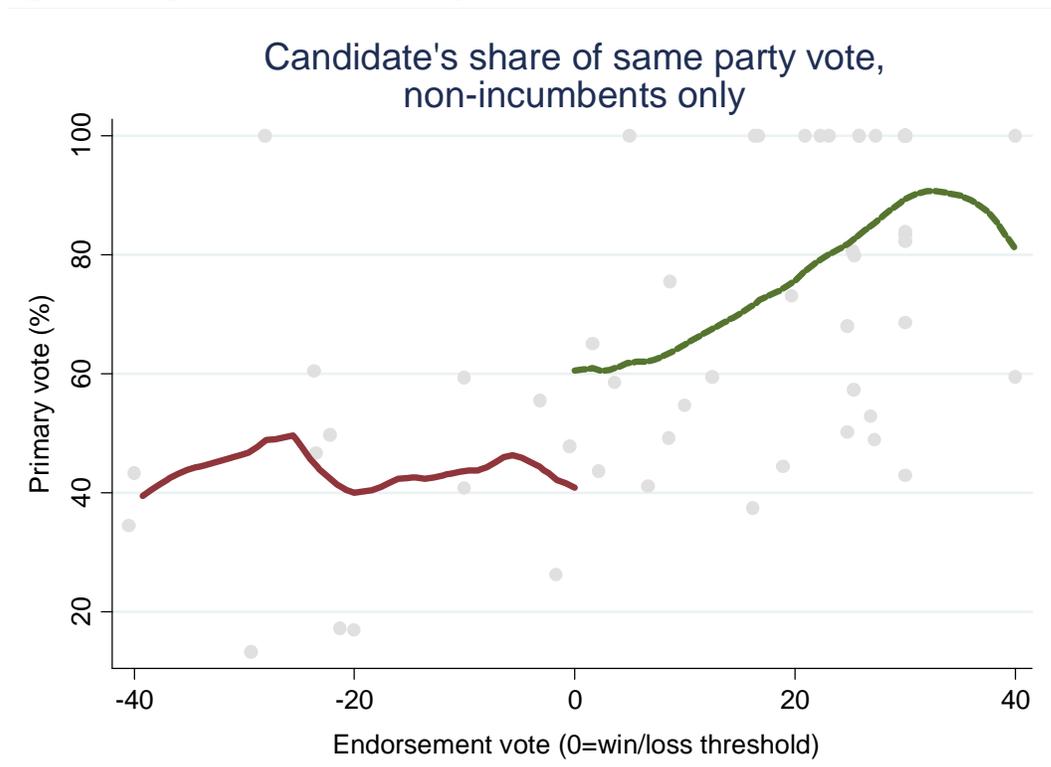
Figure 6. Fundraising Totals for Candidates Just About and Below Threshold

Money raised prior to endorsement vote

Money raised after endorsement vote
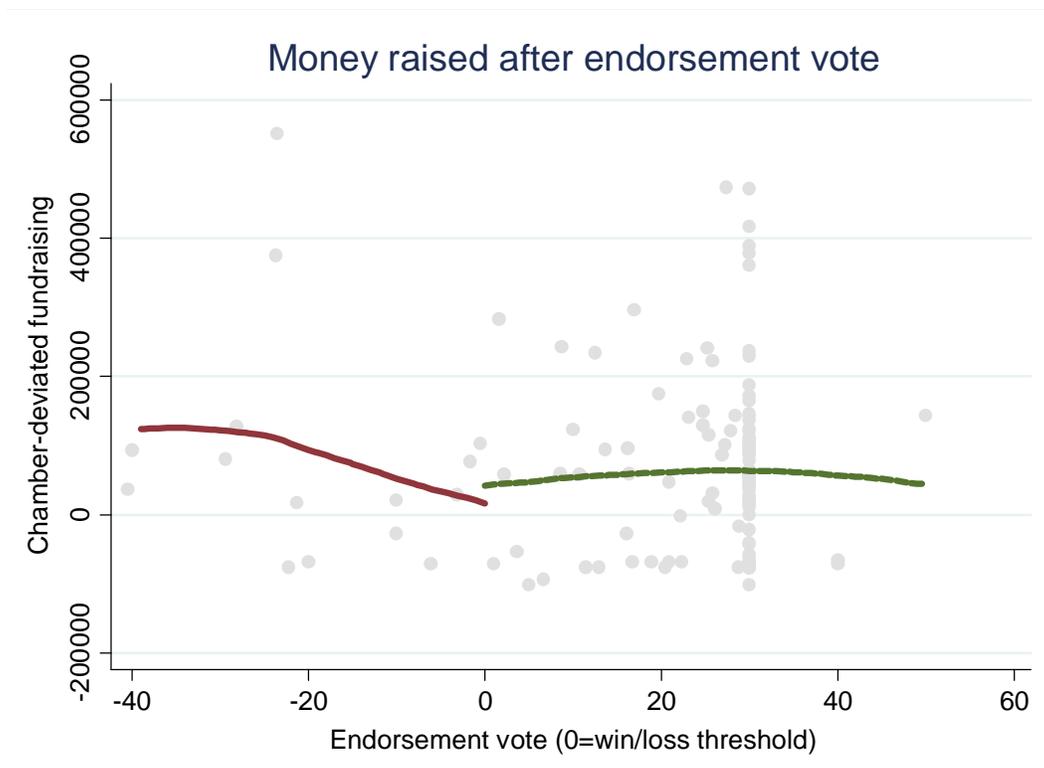
References

Ahler, Douglas J., Jack Citrin, and Gabriel S. Lenz. 2013. "Do Open Primaries Help Moderate Candidates? An Experimental Test on the 2012 California Primary." University of California.

Arceneaux, Kevin, and Robin Kolodny. 2009. Educating the Least Informed: Group Endorsements in a Grassroots Campaign. *American Journal of Political Science* 53 (4): 755-770.

Bawn, Kathleen, Marty Cohen, David Karol, Seth Masket, Hans Noel, and John Zaller. 2012. A Theory of Political Parties: Groups, Policy Demands and Nominations in American Politics. *Perspectives on Politics* 10 (3): 571-597.

Boudreau, Cheryl. 2009. Closing the Gap: When Do Cues Eliminate Differences Between Sophisticated and Unsophisticated Citizens? *Journal of Politics* 71 (3): 964-976.

Brader, Ted A., and Joshua A. Tucker. 2009. What's Left Behind When the Party's Over: Survey Experiments on the Effects of Partisan Cues in Putin's Russia. *Politics and Policy* 37 (4): 843-868.

Brader, Ted A., Joshua A. Tucker, and Dominick Duell. 2012. Which Parties Can Lead Opinion? Experimental Evidence on Partisan Cue Taking in Multiparty Democracies. *Comparative Political Studies* Online prepublication.

California Democratic Party v. Jones. 2000. 530 U.S. 567 530 U.S. 567. (U.S. Supreme Court).

Cattaneo, Matias D., Brigham Frandsen, and Rocio Titiunik. 2013. "Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate." University of Michigan.

Caughey, Devin, and Jasjeet S. Sekhon. 2011. Elections and the Regression Discontinuity Design: Lessons from Close U.S. House Races, 1942–2008. *Political Analysis* 19: 385-408.

Cohen, Marty, David Karol, Hans Noel, and John Zaller. 2008. *The party decides : presidential nominations before and after reform.* of *Chicago studies in American politics*. Chicago: University of Chicago Press.

Dominguez, Casey B. K. 2011. Does the Party Matter? Endorsements in Congressional Primaries. *Political Research Quarterly* 64 (3): 534-544.

Erikson, Robert S. 1976. The Influence of Newspaper Endorsements in Presidential Elections: The Case of 1964. *American Journal of Political Science* 20 (2): 207-233.

Fisher, Ronald A. 1935. *The Design of Experiments*. London: Oliver and Boyd.

Fiske, Susan. 1986. "Schema-Based Versus Piecemeal Politics: A Patchwork Quilt, but Not a Blanket, of Evidence." In *Political Cognition*, eds. Richard R. Lau and David O. Sears. Hillsdale, NJ: Lawrence Erlbaum.

Gimpel, James G. 1998. Packing heat at the polls: Gun ownership, interest group endorsements, and voting behavior in gubernatorial elections. *Social Science Quarterly* 79 (3): 634-648.

Gregg, James E. 1965. Newspaper Editorial Endorsements and California Elections, 1948-1962. *Journalism Quarterly* 42 (Autumn): 532-539.

Hooper, Michael. 1969. Party and Newspaper Endorsements as Predictors of Voter Choice. *Journalism Quarterly* 46 (Summer): 302-305.

Imbens, Guido, and Karthik Kalyanaraman. 2009. "Optimal Bandwidth Choice for the Regression Discontinuity Estimator." 14726. National Bureau of Economic Research.

Iyengar, Shanto, Daniel H. Lowenstein, and Seth Masket. 2001. The Stealth Campaign: Experimental Studies of Slate Mail in California. *Journal of Law and Politics* 17: 295-332.

Jacobson, Gary C., and Samuel Kernell. 1981. *Strategy and choice in congressional elections*. New Haven: Yale University Press.

Kahneman, Daniel, Paul Slovic, and Amos Tversky. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.

Keele, Luke, Corrine McConnaughy, and Ismail White. 2012. Strengthening the Experimenter's Toolbox: Statistical Estimation of Internal Validity. *American Journal of Political Science* 56 (2): 484-499.

Ladd, Jonathan McDonald, and Gabriel S. Lenz. 2009. Exploiting a Rare Communication Shift to Document the Persuasive Power of the News Media. *American Journal of Political Science* 53 (2): 394-410.

Lau, Richard R., and David P. Redlawsk. 2001. Advantages and Disadvantages of Cognitive Heuristics in Political Decision Making. *American Journal of Political Science* 45 (October): 951-971.

Lee, David S., and Thomas Lemieux. 2009. "Regression Discontinuity Designs in Economics." 14723. National Bureau of Economic Research.

Lupia, Arthur. 1994. Shortcuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections. *American Political Science Review* 88 (1): 63-76.

Lupia, Arthur, and Mathew D. McCubbins. 1998. *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* New York: Cambridge University Press.

Masket, Seth E. 2007. It Takes an Outsider: Extralegislative Organization and Partisanship in the California Assembly, 1849-2006. *American Journal of Political Science* 51 (3): 482-497.

Masket, Seth E. 2009. *No Middle Ground: How Informal Party Organizations Control Nominations and Polarize Legislatures*. Ann Arbor: University of Michigan Press.

McClenghan, Jack Sean. 1973. Effects of Endorsements in Texas Local Elections. *Journalism and Mass Communications Quarterly* 50: 363-366.

McClosky, Herbert, John Zaller, and Dennis Chong. 1985. "Social Learning and the Acquisitions of Political Norms." In *The American Ethos: Public Attitudes Toward Capitalism and Democracy*, eds. Herbert McClosky and John Zaller. Cambridge, MA: Harvard University Press.

McCrary, Justin. 2008. Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. *Journal of Econometrics* 142: 698-714.

McDermott, Monika L. 2006. Not for Members Only: Group Endorsements as Electoral Information Cues. *Political Research Quarterly* 59 (2): 249-257.
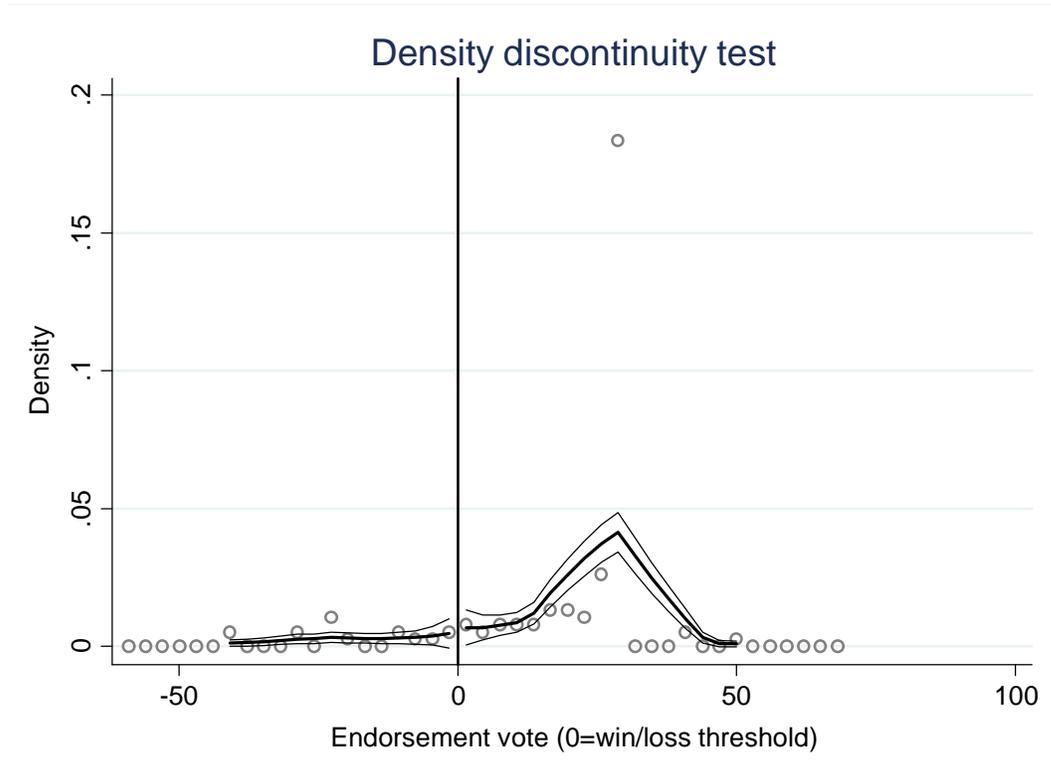
McGhee, Eric, and Daniel Krimm. 2012. *Test-Driving California's Election Reforms*. Public Policy Institute of California http://www.ppic.org/main/publication_quick.asp?i=1032.

Nichols, Austin. 2007. Causal inference with observational data. *The Stata Journal* 7 (4): 507-541.

Nichols, Austin. 2011. rd 2.0: Revised Stata module for regression discontinuity estimation. *http://ideas.repec.org/c/boc/bocode/s456888.html*.

Nisbett, Richard, and Lee Ross. 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs: Prentice-Hall.

Shor, Boris, and Nolan McCarty. 2011. The Ideological Mapping of American Legislatures. *American Political Science Review* 105 (3): 530-551.

Sniderman, Paul. 1996. Innovations in Experimental Design in Attitude Surveys. *Annual Review of Sociology* 22: 377-399.

Steger, Wayne. 2007. Who Wins Presidential Nominations and Why: An Updated Forecast of the Presidential Primary Vote. *Presidential Research Quarterly* 60: 91-97.

Stein, Lana, and Arnold Fleischmann. 1987. NEWSPAPER AND BUSINESS ENDORSEMENTS IN MUNICIPAL ELECTIONS: A Test of the Conventional Wisdom. *Journal of Urban Affairs* 9 (4): 325-336.

Stone, Walter J., Ronald B. Rapoport, and Alan I. Abramowitz. 1992. Candidate Support in Presidential Nomination Campaigns: The Case of Iowa in 1984. *Journal of Politics* 54 (4): 1074-1097.

Thistlethwaite, Donald L., and Donald T. Campbell. 1960. Regression-Discontinuity Analysis: An Alternative to the Ex-Post Facto Experiment. *Journal of Educational Psychology* 51: 309-317.

Tomz, Michael, Jason Wittenberg, and Gary King. 2003. CLARIFY: Software for Interpreting and Presenting Statistical Results. *Journal of Statistical Software* Copy at http://j.mp/k3k0rx.

Van Oot, Torey. 2011. "Republican Chairman Pushing for New Nominating Process," *CapitolAlert* (February 18). http://blogs.sacbee.com/capitolalertlatest/2011/02/crp-chair-ron-nehring.html (Accessed February 18).

Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.

Figure A1. Example of an Actual Ballot from June 5, 2012 Primary

Figure A2. Density Discontinuity Test

Density discontinuity test

---

[1] http://www.cadem.org/admin/miscdocs/files/CDP-BY-LAWS.pdf, article VIII section 2.

[2] Just before the 2012 election, 43.7% of California's registered voters belonged to the Democratic Party, 29.4% to the Republican Party, 6% were registered with a minor party and 20.9% had "no party preference." These figures are taken from the California Secretary of State's report of registration as of October 22, 2012, accessed at http://www.sos.ca.gov/elections/ror/ror-pages/15day-general-12/hist-reg-stats1.pdf in February 2013.

[3] We are grateful to Dennis Raj of the California Democratic Party for supplying us with the exact record of voters in pre-endorsement caucuses and at the state party convention, and note that these totals closely matched the figures that we observed at the convention.

[4] This logit model, estimated in Stata 11.0 with first differences simulated by CLARIFY holds constant a respondent's party identification, ideological self-placement, age, ethnicity, and indicators of whether or not she is college-educated, married, attends church frequently, and watches news frequently.

[5] This variable is often referred to as the "assignment" or "running" variable as well.

[6] We use software written for Stata by Austin Nichols to calculate this bandwidth.

[7] To calculate this density test, we use software written for Stata by Brian Kovak and Justin McCrary. The software can be downloaded at http://emlab.berkeley.edu/~jmccrary/DCdensity/.

[8] Since the victory margin for a pre-endorsement conference vote was 70 percent, candidates who received unanimous support at that stage won by a 30 point margin and had no need for a follow-up vote at the state convention. In three cases, a candidate did not receive the 70 percent margin at the pre-endorsement conference but the opposition evaporated by the convention and the final vote was unanimous. Since the margin of victory at the state level was lower (50 percent for incumbents; 60 percent for other candidates) the margin of victory was larger in those cases.

[9] The state that we study also features a split between "traditional" and "business" Democrats, but again we argue that this is not a peculiar feature. Factions within parties are nothing new in American politics, with many state Republican parties now split between Tea Party and moderate wings, while Democrats throughout much of the 20th century split over the issue of civil rights.