Supporting Information for "Partisanship as Cause, Not Consequence, of Participation"

August 18, 2021

Contents

List of Figures

List of Tables

FIGURE A.1. Partisanship in Chile

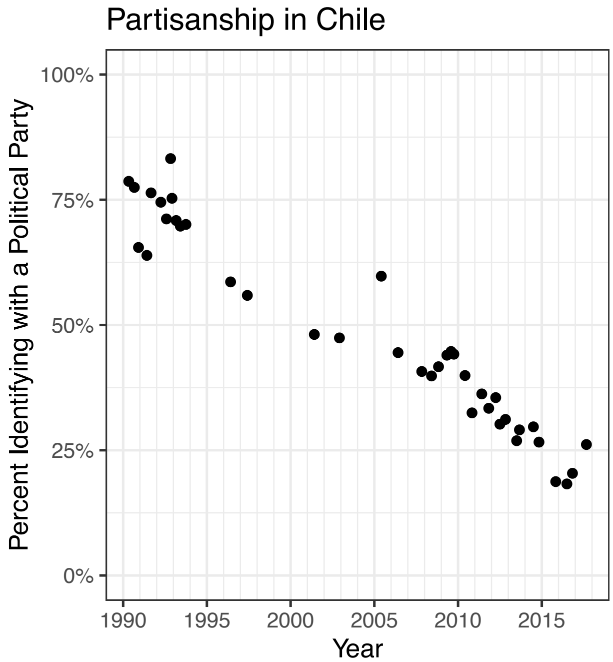

Note: Points indicate the proportion of respondents identifying with any political party in regular public opinion polls conducted by CEP (Centro de Estudios Públicos).

TABLE A.1. Partisanship and Turnout

	Dependent variable:	
	Voted in 2013 Election	
	OLS	*logistic*
	(1)	(2)
Party ID	0.175***	0.889***
	(0.028)	(0.149)
Constant	0.632	0.539
	(0.015)	(0.066)
Observations	1,350	1,350

Note: $^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$

Data: CEP November 2014 (self-reported turnout)

1 Survey Experiment

Participants in the survey experiment were recruited by Qualtrics. Surveys were conducted between August 23 and September 5, 2019. Compensation for all respondents was based on the rate for a survey of 10–15 minutes, even if they selected the 2-minute survey. Pilot results revealed that many survey respondents preferred to complete a longer survey, even in the baseline condition. While the design accounts for this (by including the baseline condition and measuring the outcome as a difference in means), I added an additional screening question to sort respondents into two groups (those who prefer long surveys and those who prefer short surveys) for the sake of optimizing the experiment's power. The only respondents included in the experiment analyzed in this paper are those who (1) indicated that they identify with a political party and (2) indicated that, all else equal, they prefer shorter surveys.

Survey length preference was measured with the following question:

> Suppose that you had the option to choose which survey to complete: a 2-minute survey or a 10-minute survey. You would receive the same payment for either survey. Which would you prefer to complete?

This question was posed at the beginning of the survey, and participants responded to a series of 4–6 demographic and political questions after this question but before treatment (to distance this question from the treatment). Among those who qualified on all other measures, 52% (635 of 1214) indicated that they prefer shorter surveys and were therefore included in the experiment. Among those 635, a total of 431 (selected at random) were included in this experiment (the other 204 were screened into a separate survey that is not part of this paper).

Supplemental Analysis

In the supplemental analysis, I compare covariates among expressive versus non-expressive respondents. Expressive respondents are those who perceive the long survey as a cost, but are willing to pay that cost to express their PID. Non-expressive respondents are those who perceive the long survey as a cost, and are unwilling to pay that cost to express their PID. These groups are defined in terms of potential outcomes. Expressive respondents will select the long survey in an expressive treatment ($y(1) = 1$) and the short survey in the baseline treatment ($y(0) = 0$). Non-expressive respondents will always select the short survey, regardless of treatment ($y(0) = 0$ and $y(1) = 0$). A third group exists: those who do not perceive the long survey as a cost ($y(0) = 1$ and $y(1) = 1$). We will call these types "survey lovers."

I assume that we do not have "defiers": respondents with $y(1) = 0$ and $y(0) = 1$. This would be an individual who prefers the long survey. In the expressive conditions, they are given the choice between completing their preferred survey (long) while getting the opportunity to express their identification, or completing their less-preferred survey (short) and not having the opportunity to express their identification. This individual would choose the latter option: completing their less-preferred survey and forgoing the opportunity to express their identification. I assume that respondents are not of this type. All respondents were screened to ensure that they identify with a political party prior to treatment. If a respondent did not want to share her party identification, she would likely respond "no" when asked if she identifies with a party (as "which party" is the natural follow-up question). Moreover, respondents could always select the "other" option if they really wished to avoid sharing the particular party with which they identify.

We never observe both potential outcomes for any individual, but we can estimate the group means for each of the three types (assuming we have no defiers). In the baseline condition, all respondents who select the long survey ($y(0) = 1$) are survey lovers. In the expressive condition, all respondents who select the short survey ($y(1) = 0$) are non-expressive respondents. The remaining groups that we observe empirically provide weighted averages for two types: the $y(0) = 0$ group is a mix of expressive respondents and non-expressive respondents; the $y(1) = 1$ group is a mix of expressive respondents and survey lovers. The $y(1) = 0$ and $y(0) = 1$ groups provide estimates of the sample proportions of non-expressive respondents and survey lovers, as well as group means for any covariate. Using these estimates, we can algebraically solve for the proportion of expressive respondents, and the group mean for any covariate among expressive respondents.

In the absence of the no-defiers assumption, these quantities are unidentifiable. We could, alternatively, compare those who opt in to the long survey versus those who select the short survey, within the expressive treatment. While this does not require the no-defiers assumption, it also sacrifices the main benefit of the experimental design: it simply compares respondents who opt to express their identification and complete the long survey with those who refrain from expressing their partisan identification and complete the short survey. But it does not distinguish between those who express their identification at a cost (expressive respondents) and those who enjoy long surveys. This approach just allows us to measure how

FIGURE A.2. Reasons for Voting

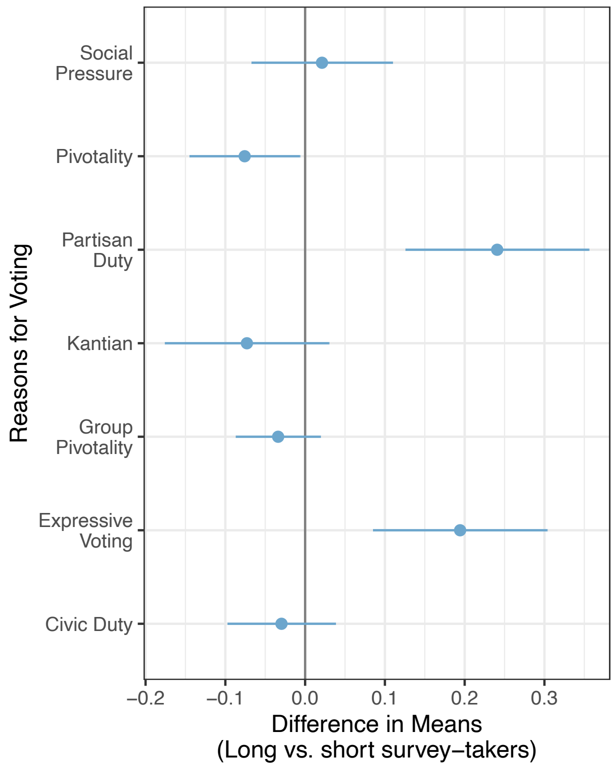

Note: Estimates come from bivariate regressions (reason for voting regressed on survey selection).

well the outcome measure (completing a long survey) correlates with other characteristics, such as self-reported turnout. I present the results for these calculations below, which are broadly consistent with the previously computed differences between expressive and non-expressive respondents.

With respect to self-reported turnout, those who opt in to the longer survey (in the expressive treatment) are 20 points more likely to report always voting ($p < 0.01$). Figure A.2 illustrates the difference in means estimates for each reason for voting, with 95% confidence intervals.

TABLE A.2. Reasons for Voting (Experiment)

Label	Reason
Expressive Voting	To express my support for my party
Partisan Duty	To contribute to my party's electoral success
Civic Duty	To fulfill my civic duty
Pivotality	My vote could change the outcome of the election
Group Pivotality	Together, my vote and the votes of people like me could change the outcome of the election
Social Pressure	If I didn't vote, people would judge me

Note: Respondents were asked to indicate the extent to which they agreed/disagreed with each reason when thinking about their own decision to vote. The sample was limited to those who indicated that they vote at least occasionally.

2 Regression Discontinuity

The regression discontinuity was estimated using survey data from the Centro de Estudios Públicos (CEP) survey project. Table A.3 lists the survey waves that were used in the RD calculations, along with the dates of data collection and the number of observations. Table A.4 lists the date and type of each election held after the 1988 plebiscite (through the 2010 election). Each respondent was surveyed after 12–15 post-plebiscite elections had occurred (with an average of 13.3 elections).

TABLE A.3. Survey Data

Wave	Dates	Observations
52	Jun–Jul 2006	1417
54	Dec 2006	1438
55	Jun 2007	1426
56	Nov–Dec 2007	1397
58	Nov–Dec 2008	1417
59	May–Jun 2009	1069
60	Aug 2009	1438
61	Oct 2009	1428
62	Jun–Jul 2010	1417
63	Nov–Dec 2010	1322
64	Jun–Jul 2011	1446
65	Nov–Dec 2011	1473
Total		16,688

A McCrary sorting test revealed no apparent sorting of the running variable (recorded birth date) around the cutpoint for the RD. The test was conducted using the `DCdensity` function from the `rdd` package in R. (See Fig. A.3.) The test yielded $\theta = -0.046$, $\sigma = .051$, $p = 0.372$.

TABLE A.4. Elections (post-plebiscite)

Date	Type
Dec 14, 1989	Presidential
Jun 28, 1992	Municipal
Dec 11, 1993	Presidential
Oct 27, 1996	Municipal
Dec 11, 1997	Parliamentary
Dec 12, 1999	Presidential (first round)
Jan 16, 2000	Presidential (second round)
Oct 29, 2000	Municipal
Dec 16, 2001	Parliamentary
Oct 31, 2004	Municipal
Dec 11, 2005	Presidential (first round)
Jan 15, 2006	Presidential (second round)
Oct 26, 2008	Municipal
Dec 13, 2009	Presidential (first round)
Jan 17, 2010	Presidential (second round)

2.1 Estimating Compliance

Using the self-reported registration data, we can estimate the proportion of compliers in the sample. But we know that respondents over-report electoral participation, due to social desirability bias. Thus, Figure A.4 presents three separate estimates of the discontinuity in registration rates. The first graph simply uses the self-reported registration status of respondents. The second and third graphs, however, use official registration numbers to estimate over-reporting and adjust the data accordingly.

When adjusting for over-reporting, I limit the sample to surveys conducted after the 2009-2010 election (n=5511). We know the actual registration rate among the entire population in 2010, and among those eligible for the plebiscite in 1988. We also know the proportion who were eligible for the plebiscite. Using these three pieces of information, we can calculate the over-reporting rate among citizens eligible for the plebiscite, and among citizens ineligible for the plebiscite.

Ninety-four percent of those eligible for the plebiscite report that they were registered to vote in the 2010 election, but only 92% of the population registered for the plebiscite. I assume that if someone was eligible for the plebiscite and chose not to register in 1988, then they did not register in later years. Previous empirical studies of registration in Chile have shown that this is a reasonable assumption (Corvalan and Cox 2013). It is also a conservative assumption: it uses the lower-bound on registration within the treatment group (where the registration rate within the treatment group represents the sum of compliers and always-takers).

Across the full sample, 74% of respondents report registering to vote, but overall registration for the 2009–2010 elections was only 68%. Sixty-two percent of respondents were

FIGURE A.3. McCrary Sorting Test

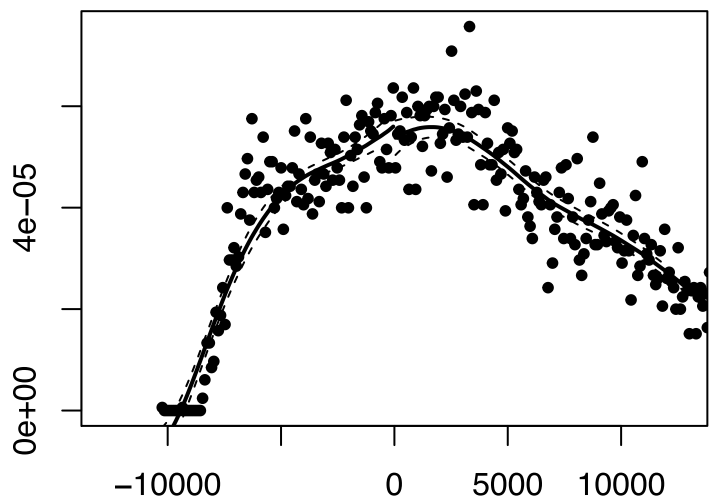

Note: The x-axis indicates the number of days before or after the registration cutoff that an individual was born.

eligible for the plebiscite. If plebiscite-eligible voters registered at a rate of 92%, then plebiscite-ineligible voters must have registered at a rate of 29% to yield an aggregate registration rate of 68%. Those who were ineligible for the plebiscite report that they registered at a much higher rate: 42%.

We can adjust the self-reported registration rates to account for this over-reporting in either of two ways. First, we can apply a proportional adjustment. If 42% of plebiscite-ineligible citizens report registering to vote when only 29% did, then there is a 31% chance that any particular self-reported registrant is actually un-registered. So to adjust the registration rates, I randomly assign 31% of self-reported registrants (in the plebiscite-ineligible group) to un-registered status. Similarly, 2.1% of self-reported registrants in the plebiscite-eligible group are assigned to un-registered status (0.02/0.94).

Second, we can apply a flat adjustment: we take the trend line from the self-reported registration data and simply shift it down by the over-reporting rate. For plebiscite-eligible voters, this rate is 2% (0.94-0.92). For plebiscite-ineligible voters, this rate is 13% (0.42-0.29).

Using the raw, unadjusted numbers, we get the most conservative estimate of the discon-

FIGURE A.4. Discontinuity in Registration

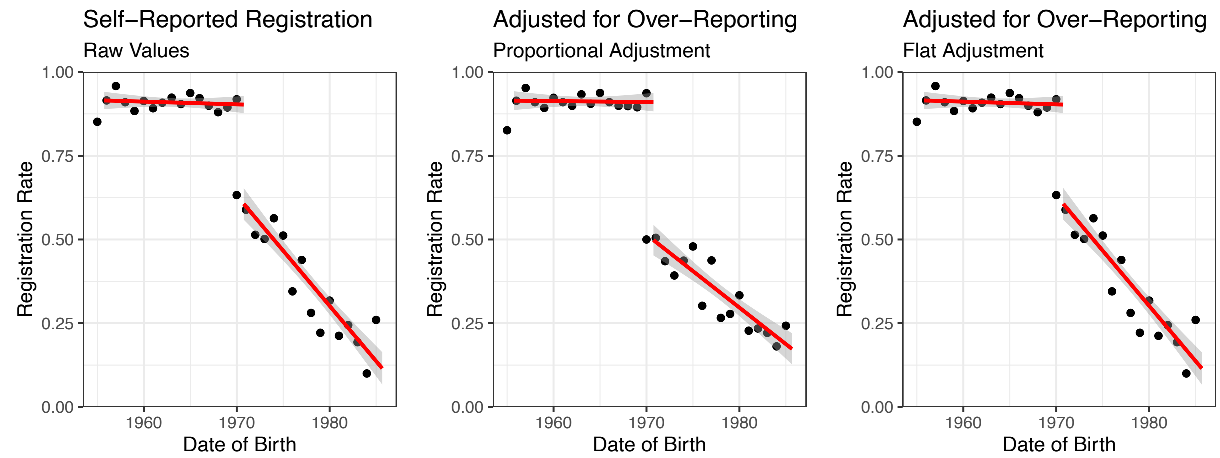

Note: Effect of plebiscite eligibility on registration for the 2009–2010 presidential election. Surveys were conducted from June 2010–December 2011 (n=5511). The left-most graph presents self-reported registration data, with no adjustments. The other graphs adjust the self-reported registration rates, using official data to estimate social desirability bias.

tinuity: approximately 16 percentage points. Applying the proportional adjustment yields an estimated effect of 40 percentage points. And applying the flat adjustment yields an estimated effect of 28 percentage points. Recall that these figures are estimates of the proportion of compliers in the sample—those citizens who would register to vote if eligible for the plebiscite, but would not register otherwise.

2.2 Additional Robustness Tests

Figures A.5 and A.6 replicate the main analysis (Fig. 4) with alternative specifications. Figure A.5 uses the same model as the main analysis, but introduces fixed effects for the survey wave. Figure A.6 uses the same model as the main analysis, but uses identification with a coalition as the dependent variable (in place of identification with a party).

Figure A.7 illustrates the difference in means for all bandwidths up to 8 years. Across all of these bandwidths, we never observe a positive treatment effect. When we use the difference-in-means comparison, rather than the RD setup, we must make assumptions about similarity in relevant covariates across the sample. Treatment assignment (a birthdate before or after the plebiscite cutoff) must be orthogonal to other factors that affect partisan identification. Within treatment groups, we do not observe any significant relationship between age and partisanship (see Fig. 3). Nonetheless, the necessary assumptions become stronger as the bandwidth increases, so I focus on smaller bandwidths here than the RD optimal bandwidths.

At the smallest bandwidths, we observe noisy estimates due to small sample sizes, with a handful of large negative (generally insignificant) effects. Of the 2917 bandwidths tested,

FIGURE A.5. RD Coefficient Plot: Survey Wave Control

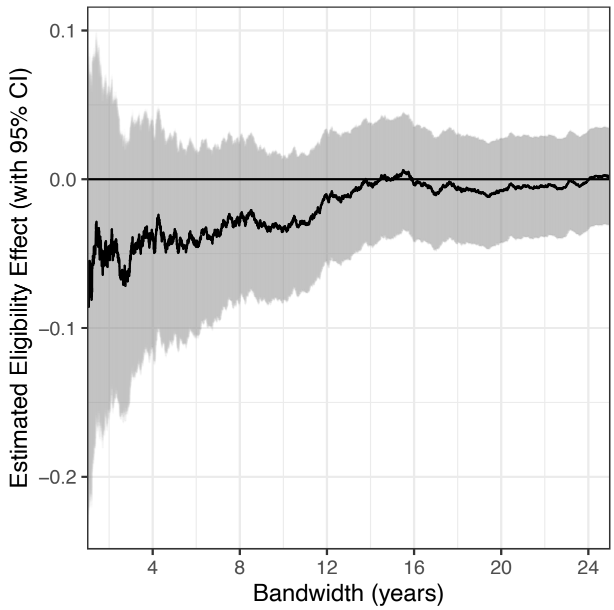

Note: Coefficient plot for each estimate of Eq. 1, using bandwidths from 1 to 25 years (increasing in intervals of one day). For each bandwidth, I plot the point estimate of β_1 with a 95% confidence interval. In this model, I control for survey wave (using a dummy for each different survey wave).

16 produce statistically significant effects (all negative). These occur at bandwidths from 27 to 49 days (with 51 to 102 observations). The effect quickly trends towards zero as the sample size increases. This test is well-powered to detect ITT effects smaller than 0.04 at about a 5-year bandwidth (see Fig. A.9 for power estimates across bandwidths).

FIGURE A.6. RD Coefficient Plot: Coalition Identification

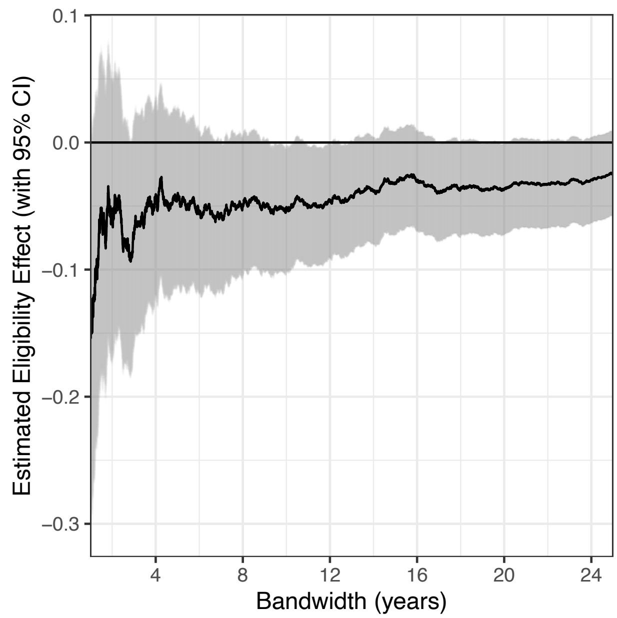

Note: Coefficient plot for each estimate of Eq. 1, using bandwidths from 1 to 25 years (increasing in intervals of one day). For each bandwidth, I plot the point estimate of β_1 with a 95% confidence interval. In this model, I use coalition identification as the dependent variable (instead of party identification).

FIGURE A.7. Effect of Eligibility on PID

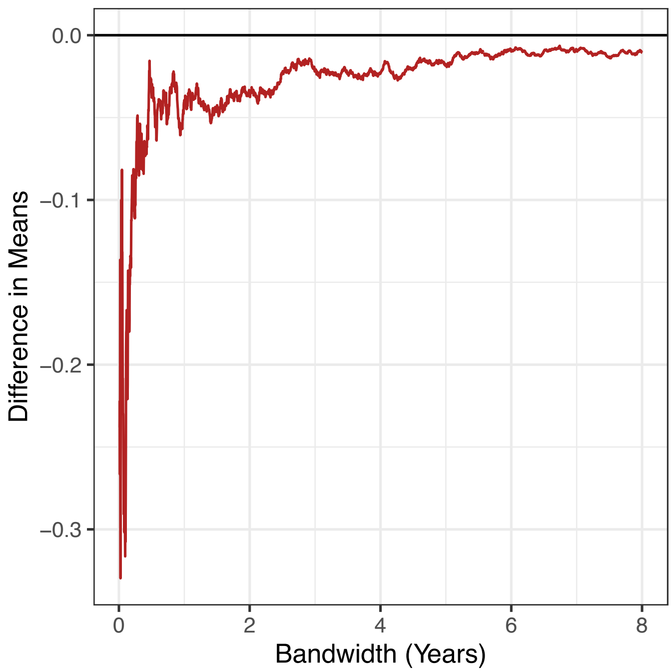

Note: Effect of plebiscite eligibility on partisan identification using the raw difference in means. All possible bandwidths (increasing by one-day intervals) up to 8 years are presented here.

2.3 Power Analysis

Figures A.8 and A.9 illustrate the statistical power of the regression discontinuity test using local linear regression and the difference in means. The curve plots the minimum β_1 effect size that the test can detect at 80% power for each bandwidth ($p < 0.05$, one-tailed tests). Each curve was generated through simulations.

FIGURE A.8. Statistical Power: RD Test

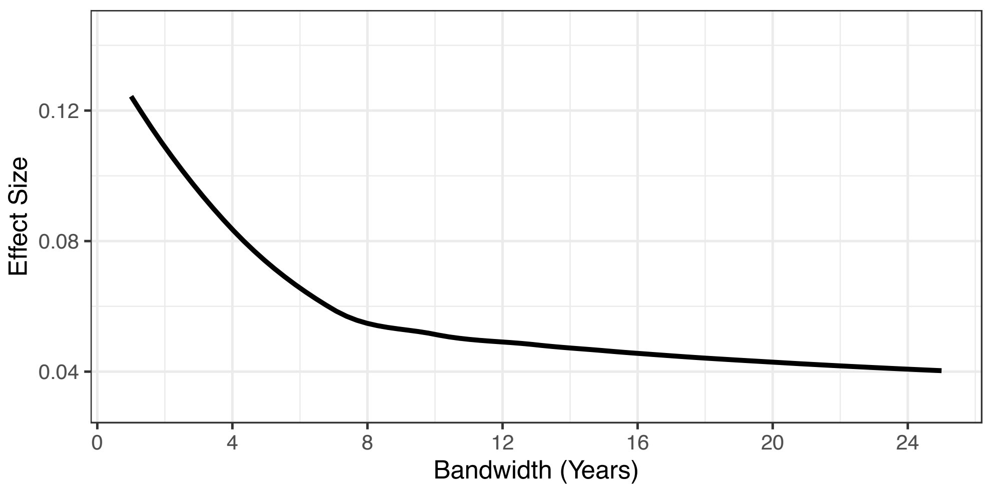

Note: 80% power at each bandwidth for the regression discontinuity estimated in Fig. 4.

FIGURE A.9. Statistical Power: Difference in Means

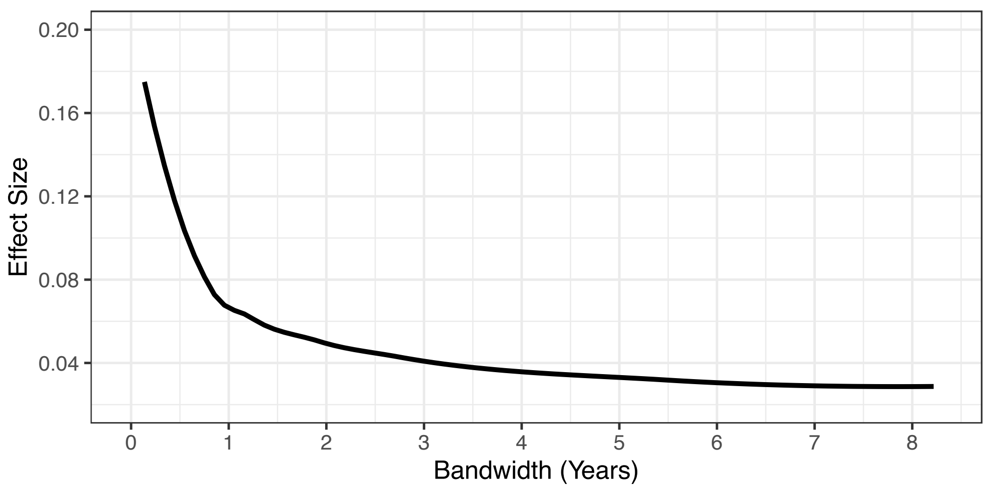

Note: 80% power at each bandwidth for the difference in means estimated in Fig. A.7.

Recall that β_1 (the effect size referenced in these power calculations) is the ITT, measuring the effect of plebiscite eligibility on partisanship. The LATE interpretation of β_1

— the effect of *voting* on partisanship — depends on assumptions about the model of the cumulative effects of voting in multiple elections. The LATE interpretation presented in the text assumes a simple model where a non-partisan has a certain probability of adopting a partisan identification each time she votes as a non-partisan. In other words, when she votes in her first election, she adopts a partisan identity with probability π. If she adopts a partisan identity after the first election, she remains a partisan after voting in the second election. If she does not adopt a partisan identity after the first election, she again has probability π of adopting a partisan identity once she votes in the second election. Thus, after voting in two elections, her probability of being a partisan is $\pi + \pi(1 - \pi)$. Alternative models might consider a waning effect: perhaps if a voter participates in many elections without developing a partisan identity, she becomes very unlikely to develop one from voting in future elections; or a more cumulative process: perhaps voting in a single election rarely leads to partisanship, but voting in three elections has a big effect, with the repetition of experience generating a sort of tipping point. Any such alternative model will affect the LATE interpretation of these ITT power calculations.

3 Observational Survey

The survey was conducted online from March 21–April 10, 2019 on a sample collected through Qualtrics. The sample included 598 total respondents.

Initially, I determined quotas for age, education, gender, and region, using the most recent census data. Partway through fielding the survey, Qualtrics determined that they did not have sufficient access to older respondents and respondents with low education. We relaxed and eventually removed the quotas to fill the sample, which skews younger and more educated than the Chilean population at large. Table A.5 presents detailed information on the sample demographics in comparison to the census data.

TABLE A.5. Sample Demographics

Category	Census data	Sample count	Sample proportion
Age			
18–24	14.25%	105	18%
25–34	20.79%	163	27%
35–44	18.07%	133	22%
45–54	17.62%	115	19%
55–64	14.23%	47	8%
65+	15.05%	35	6%
Gender			
Male	48.95%	251	42%
Female	51.05%	347	58%
Education			
Less than secondary	37.78%	10	2%
Complete secondary	31.97%	89	15%
Incomplete vocational	1.58%	52	9%
Complete vocational	7.48%	116	19%
Incomplete university	6.00%	99	17%
Complete university	15.18%	232	39%

With any online survey, it is not possible to collect a traditional probability sample of the general population. We cannot state with certainty how the survey population differs from the population at large, because there may be other unobserved characteristics on which the survey respondents differ from the broader population. However, using well-studied characteristics like age and education, we can draw inferences about how the sample might differ.

As illustrated with the CEP data used in the regression discontinuity, age does not bear any strong relationship to partisanship in Chile. It does correlate with voting experience and propensity to turn out to vote (now that voting is voluntary), with younger people being less likely to turn out to vote. A younger sample, then, likely has a lower electoral participation rate than the population average.

The education skew pushes in the opposite direction on participation: turnout rates rise with education (see, e.g., Corvalan and Cox 2013). Education is also positively correlated with partisan identification. Figure A.10 illustrates the relationship between education and partisanship according to wave 82 of the CEP surveys, conducted in October and November 2018.

FIGURE A.10. Education and Partisan Identification

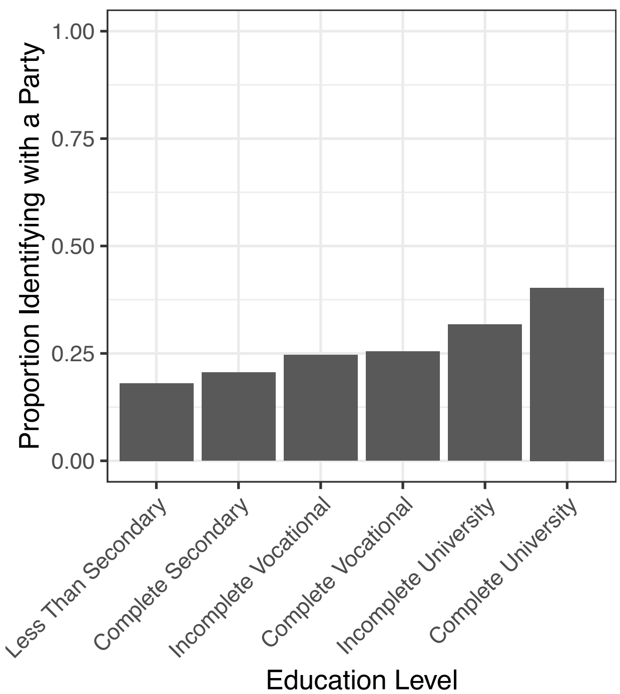

Note: Estimates come from the CEP wave 82 survey, conducted in October-November 2018.

In line with these patterns, respondents indicate higher than average rates of partisanship: 55% of respondents in my survey indicated a party identification in the first question (that is, the standard question asking which party the respondent feels closest to), compared to 23.3% in the CEP wave 82 survey. Although the sample has a greater number of partisans than a probability sample would yield, the implications for representativeness among partisans are unclear — conditioning on having a partisan identification, it is not clear whether partisanship is stronger or weaker among partisans in my survey, compared to the broader population of partisans. Establishing such claims of broader representativeness would require implementing the more detailed partisanship questions used here within a survey conducted on a probability sample of the general population.

Figure A.11 presents the distribution of scores on the partisan strength scale. All respondents were asked which party they identified with most strongly. Those who indicated that they do not identify with any political party saw a follow-up question, asking which party they feel a little closer to. Some respondents (169) still refused to select a party, so they did not see the follow-up questions for the partisan scale.

FIGURE A.11. Strength of partisan identities.

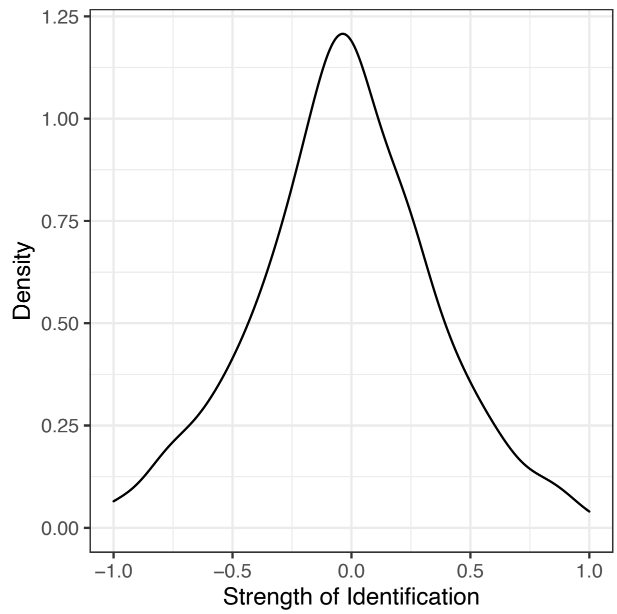

Note: Distribution of partisan strength scores (n=429).

The partisan scale was constructed by adding together the responses from eight questions, and re-scaling to a 0 to 1 scale. For each question, response options included “disagree strongly,” “disagree somewhat,” “agree somewhat,” and “agree strongly.”

Question wording: You indicated earlier that you identify most strongly with [PARTY]. Please indicate the extent to which you agree with the following statements, thinking about [PARTY].

- When I speak about this party, I usually say ”we” instead of ”they.”
- I am interested in what other people think about this party.
- When people criticize this party, it feels like a personal insult.
- I have a lot in common with other supporters of this party.
- If this party does badly in opinion polls, my day is ruined.
- When I meet someone who supports this party, I feel connected with this person.
- When I speak about this party, I refer to them as ”my party.”
- When people praise this party, it makes me feel good.

TABLE A.6. Reasons for Voting (Observational)

Label	Reason
Expressive Voting	To express my support for my party
Partisan Duty	To contribute to my party's electoral success
Civic Duty	To fulfill my civic duty
Pivotality	My vote could change the outcome of the election
Group Pivotality	Together, my vote and the votes of people like me could change the outcome of the election
Social Pressure	If I didn't vote, people would judge me
NA	I don't vote

Note: Respondents were asked to indicate the most important reason why they vote, from the options listed under "Description."