

RESEARCH NOTE

Early voting experiences and habit formation

Elias Dinas , Vicente Valentim* , Nikolaj Broberg and Mark N. Franklin 

European University Institute, Fiesole, Italy

*Corresponding author. Email: vicente.valentim@eui.eu

(Received 16 July 2020; revised 9 March 2021; accepted 7 June 2021; first published online 14 March 2023)

Abstract

Research has shown that first-time voting experiences affect subsequent voting behavior, with salient elections boosting subsequent turnout and non-salient ones suppressing it. We challenge this view. Following research on the context-dependent nature of habit formation, we argue that all elections should affect subsequent turnout in elections of the same type. Comparing individuals that differ only in how salient their first eligible election was (Presidential or Midterm), we find support for this expectation. Individuals are more likely to vote for, and be interested in, elections of the same type as their first voting experience. Leveraging voting age laws in the US, we also show that such laws affect subsequent participation by changing the type of election individuals are first eligible for.

Keywords: Eligibility; midterm elections; presidential elections; turnout; voting habits

There is growing consensus among students of voting behavior that the electoral setting individuals encounter in their first eligible election is crucial in shaping future voting patterns (Meredith, 2009; Dinas, 2012; Coppock and Green, 2016; de Kadt, 2017). The dominant view in this literature argues that this effect is due to the resources individuals gather from their initial voting experiences, which can make it easier for them to cast a vote later on (Gerber *et al.*, 2003). An implication of this argument is that the direction of the effect should depend on the salience of the first election one is eligible to vote for. More salient elections are deemed to boost future turnout (Meredith, 2009) by generating voting habits (Plutzer, 2002; Franklin, 2004). Less salient elections are hypothesized to either leave no imprint on future voting trajectories (Meredith, 2009), or to build non-voting habits that undermine future electoral participation (Franklin and Hobolt, 2011). This conclusion can fuel arguments against the existence of low-salience elections such as the European elections in EU member states, which could create abstentionist habits in voters.

There are, however, theoretical grounds to challenge this dominant view. Habits tend to be formed by individuals recognizing patterns in the co-occurrence of a given action and the contextual setting in which such action is performed (Hall, 1969; Wood and Neal, 2007). Once cued with the similarity between the present context and a previous one, individuals tend to respond to it with the same behavior (Wood and Neal, 2007). For this reason, contextual continuity is crucial for the formation of habits.

A number of empirical studies have shown that the formation of human habits depends on the continuity of the context in which they take place. Individuals tend to perform a number of activities mostly under a similar context. For example, they tend to watch the news on television on the same channel, at the same time of the day; and they tend to purchase similar foods at the same time of the day (Ji and Wood, 2007). Once created, these habits can last even if they are inefficient (Larcom *et al.*, 2017). Conversely, disruptions to contextual continuity are likely to lead individuals to break habits (Wood *et al.*, 2005).

Similar to other types of behavior, voting behavior can also be contingent upon the context. If people develop regularities in their responses to familiar circumstances, the repeating context surrounding elections is likely to matter. Much like other instances of institutionalized collective action, similarity in the electoral context can provide individuals with the necessary cue to activate their habitual response (Aldrich *et al.*, 2011).

Empirical evidence in political science has produced findings that mirror those obtained from social psychology when looking at different types of behavior. Continuity in the context of an election strengthens the effect of past vote on future elections. Conversely, disruptions to such continuity—such as those when individuals move to a different location—hamper that effect (Green and Shachar, 2000; Denny and Doyle, 2009; Aldrich *et al.*, 2011).

We focus on another factor that should be expected to forge contextual continuity: the type of electoral battle. Different types of elections are associated with different types of campaigns and generate different electoral stimuli. As argued by de Kadt (2017), if voting creates a positive emotional state, then future opportunities to enter that state will be attractive. Extending this line of argument, we expect same-type elections to present more evident opportunities to reach this state. Early voting experiences should thus have a higher effect on turnout in subsequent elections of the same type. Voting for the first time for a high-salience election should have a higher effect on turnout in subsequent high-salience elections; voting for the first time for a low-salience election should have a higher effect on turnout in subsequent low-salience elections. This expectation comes in stark contrast with that of the salience argument, which would expect high-salience elections to have a higher effect on subsequent voting behavior in elections of any type.

Previous research enabling us to discriminate between these two mechanisms remains scarce. Existing evidence stems primarily from studies that leverage the fact that national elections in European countries are more salient than elections for the European Parliament (EP). Their results, however, remain ambiguous. Franklin and Hobolt (2011) find that eligibility to vote in an election for the European rather than the national parliament decreases turnout in subsequent elections, while Dinas and Riera (2018) show that EP versus national eligibility increases future vote for small parties. Against these results, Schulte-Cloos (2019) finds that EP eligibility increases interest in politics without suppressing the vote for mainstream parties.

The most rigorous evidence in favor of the electoral saliency hypothesis comes from Meredith (2009). Looking at California voting records, this author finds that past participation in Presidential elections induces future electoral participation in the 2004 and 2008 Presidential elections. He also finds no effects of 2002 eligibility on vote in either the 2004 or the 2006 election. Yet, as Meredith points out, this finding might not be due to the different nature of the elections but rather due to the fact that turnout decisions among those who already vote in Midterm elections are less likely to be affected by past participation. Moreover, these findings are at odds with those of Coppock and Green (2016). Using eligibility discontinuities in 17 US states, these authors find that being eligible to vote in the upstream election significantly increases voting in downstream elections. Unlike Meredith (2009), they find that the effect holds both for Presidential and Midterm elections.

To provide a more authoritative answer to this question, we rely on a research design specifically conceived to test the salience hypothesis against the contextual-continuity hypothesis. We take advantage of the fact that, in the US, high salience Presidential elections happen intertwined with low salience Midterm elections. Using all publicly available American National Election Studies (ANES), we exploit as-good-as-random variation in the type of election that respondents were first eligible to vote in. Instead of comparing individuals on the basis of their eligibility in a given election, our design compares them on the basis of the type of election for which they were first eligible. Doing so means that our design does not allow us to test the effect of eligibility *per se* (which previous literature has provided evidence on). Instead, our design is meant to evaluate the conditions under which the effect of eligibility is stronger—is it stronger for individuals first eligible to vote in a highly salient election (electoral saliency hypothesis), or

when the election type matches that of the first election respondents were eligible to vote in (contextual-continuity hypothesis)?

As a last step in the empirical analyses, we examine a policy implication of the theory. If individuals develop context-specific habits, policies that change the type of election individuals are first eligible for should have lasting consequences on their subsequent voting behavior. We test for this expectation by looking into changes to voting age legislation in the US. This law made some individuals who would otherwise be first eligible to vote for a Presidential election be first eligible for a Midterm election. Our analyses show that such change made them more likely to vote for subsequent Midterm elections, but less so for Presidential elections.

1. Data and empirical strategy

We use all ANES data covering elections from 1956 to 2018.¹ Our outcome of interest is self-reported turnout in election s_n , predicted by a dummy that denotes whether election s_n is of same type (Presidential or Midterm) as election s_1 —each respondents' first eligible election. We obtain this information by subtracting individuals' age from the year of the survey, which provides us with their year of birth. Knowing the year of birth allows us to predict with high accuracy individuals' first eligible election.² We use no information as to whether respondents actually voted in their first eligible election. Rather, we stick to the reduced form, employing type-of-election eligibility as a predictor of turnout in the year of the survey. Moreover, we omit from the analysis the elections used to classify respondents in terms of their treatment status, i.e. individuals' first pair of Presidential and Congressional elections. Additional information on the data can be found in the online Appendix.

Identification relies on two assumptions: that potential turnout rates in same-type elections would be the same for both groups had there been no difference in their type of first-eligible election (ignorability); and that type-of-first-election eligibility does not affect future turnout in same-type elections through other ways except via its effect on the type of election one encounters when coming of age (exclusion).³ The first assumption, ignorability, would be violated if individuals (or their parents) could in some way select themselves (or their children) into the type of first-eligible election, a rather unlikely scenario.

Exclusion would be violated if type-of-election eligibility alone affected turnout patterns without this effect going through the experience of the election itself. This could be the case if for example Presidential-eligibles ended up being for some reason older or more educated than Midterm-eligibles. Figure 1 illustrates the logic upon which the design is based. The alternation of black and grey stripes in the left-hand panel suggests the absence of any monotone association between year-of-birth and type-of-election eligibility. It is this alteration we aim to leverage for identification purposes. We believe that eventually belonging in one of the two groups is as-good-as-random. The right-hand side of the graph, which compares the age distribution between the two groups, does suggest a high degree of overlap. That said, we do find some imbalances in the age and level of education of respondents, which, as Table I shows, are statistically significant. To address this concern, we employ various specifications, including all or some of these variables as covariates. We also include a dummy for whether the current election is Presidential or Midterm. Finally, since survey respondents vary in age, they also vary in terms

¹The only exceptions are the Midterm elections that took place in 2010 and 2014, for which ANES does not provide a survey in the Data Center section on their website.

²Uncertainty occurs only when the respondent becomes of age on the year of the election. Because elections happen in November, we assume that these individuals already had their birthday and thus we treat them as eligible to vote in that election. We show that results remain unchanged when we relax this assumption.

³Given that the estimand we use here is the Intent-To-Treat, exclusion is, strictly speaking, not needed. Yet, we choose to discuss it because we believe it matters for our interpretation of the effects, which rests on the idea that the effects are due to the election itself.




Figure 1. Visualizing the identification strategy: comparing groups according to the type of first eligible election.

Table 1. Difference in means between Presidential and Midterm eligibles

	Presidential eligibles	Midterm eligibles	Difference of means	Standard deviation	Number of observations
Education: grade school or less	0.112	0.127	0.014***	0.324	55770
Education: high school	0.432	0.442	0.011**	0.496	55770
Education: some college	0.237	0.228	-0.009***	0.423	55770
Education: college or advanced degree	0.219	0.203	-0.015***	0.409	55770
White	0.770	0.770	0.000	0.422	59419
Age	46.333	47.085	0.752***	17.26	59006
Age ²	2439.39	2519.23	79.84***	1732.81	59006
South	0.301	0.302	0.001	0.459	56244

of how distant the outcome election, s_n , is from the treatment election, s_1 . For some people their outcome election is the third or fourth Presidential or Midterm election in which they are of age to vote, while for others it might be their 10th eligible election. To account for such heterogeneity, we include election-counter fixed effects.

Another way in which exclusion could be violated is if Presidential- or Midterm-year eligibility affects future turnout in Presidential or Midterm elections respectively via other routes than the first-eligible election itself. For example, if high-school curriculum also changes in years of Presidential or Midterm elections, it could be that effects are not due to the election itself but rather due to cohort-specific changes in school curriculum.⁴ We cannot test this assumption, but we try to gauge whether it holds by looking into several post-treatment outcomes that are not election-specific. Our treatment is based on the idea that similar contextual stimuli as those experienced in first eligible election should moderate behavioral responses. By extension, we expect that attitudinal outcomes not distinguishing between types of elections remain unaffected by our treatment. We use all such outcomes available in the cumulative ANES survey, as shown in Table 2. The means of the two groups are close to identical across all outcomes. Presidential- and Midterm-eligibles do not differ in terms of their attitudes toward government spending; whether they approve or not the President; whether they think that parties are only

⁴To be sure, if these changes are produced by the type of election on the specific year, these effects could be still attributed, even if indirectly, to the election itself.

Table 2. Differences between midterm and presidential eligibles on political outcomes

	Presidential eligibles	Midterm eligibles	Difference of means	p-value	Number of observations
Urbanism	2.120	2.110	-0.010	0.206	40014
Moved	0.916	0.913	0.003	0.238	55800
Approved	0.570	0.568	-0.001	0.779	40568
President					
parties only	0.603	0.591	-0.012	0.191	12071
interested in votes					
High	0.330	0.325	-0.005	0.220	50694
income					
Region:	0.192	0.191	-0.000	0.894	55800
northeast					
Region:	0.262	0.265	0.002	0.531	55800
north central					
Region:	0.355	0.355	-0.000	0.917	55800
south					
Region:	0.191	0.189	-0.001	0.658	55800
west					
Voting	0.103	0.104	0.001	0.815	14112
matters					
Politics	0.711	0.702	0.009	0.068	34720
complicated					
Voters	0.632	0.651	0.019	0.009	17968
have a say					
Internal	0.441	0.445	0.003	0.483	41517
efficacy					
External	0.551	0.555	0.004	0.439	40767
efficacy					
Trust	2.439	2.445	0.006	0.323	39324
government					
Government	139.9105	141.9942	2.083748	0.4186	44215
responsiveness (index)					
party wins					

interested in casting votes; or whether parties differ. They are also indistinguishable when it comes to regional residence as well as residential mobility. Finally, when looking also at the supplementary items included in the ANES (available for only a third of the total number of respondents) we find that the two groups do not differ in terms of internal or external efficacy, trust in government, opinions about whether voters have a say in politics; whether voting matters; or whether politics is too complicated. Among these 17 items, only one indicates a difference which albeit very small is significant at the 5 percent level. Figure 2 shows that when arraying the p-values obtained from this exercise against a uniform distribution, we find only small variations from the 45 degree line, which corresponds to the level of balance expected under true randomization. Without being in any way exhaustive, the evidence enhances our confidence that the effects reported here are due to type-of-election eligibility.

Thus, our baseline model takes the following form:

$$Vote_{i,s>1} = \sum_{s=2}^{15} \alpha_s E_s + \lambda_1 Same_{is} + \delta P_s + X\beta_2 + u_i \tag{1}$$

where $Same_{is}$ is coded as one if the election in question is of same type as first-eligible election and zero otherwise. E_s indexes election-counter dummies, where $s = \{2, \dots, 15\}$. Every s represents two elections, one Midterm and one Presidential, which together comprise a pair. Thus, $s = 2$ means that i is interviewed after the second election she was eligible to participate, either




Figure 2. Balance on observables between Presidential and Midterm eligibles.
 Note: The graph displays p-values corresponding to a difference-of-means test between Presidential and Midterm eligibles. The dashed diagonal line denotes the 45-degree line, which is where the p-values would lie, if they came from a uniform distribution. The blue dashed lines denote the 5 percent cutoff point and the red dashed lines denote the 10 percent level.

Presidential or Midterm, and so on, up until $s = 15$. On average, our respondents encounter their first eligible election when they are 20 years old. Thus, going up to $s = 15$ means that we trace a 60-year-long life trajectory, from 20 to approximately 80 years of age. P_s is a dummy for Presidential elections that controls for differences in turnout between these elections and midterm ones. X is a vector of control variables. In our preferred specification, we include all covariates included in Table 1.⁵ In the online Appendix we use two alternative specifications. The first includes only education and age, the only two covariates we deem necessary for identification.⁶ The second is intended to assess the robustness of these estimates by including contemporaneous sociodemographics as proxies for socioeconomic background: income, region of residence by the time of the ANES interview, and a trichotomy distinguishing respondents between those residing in central cities, suburban cities, or rural towns (*Urbanism*).⁷ Conditional on these observables, we assume as-good-as-random assignment of type of first eligible election. If this is the case, λ_1 recovers the effect of same-type election on probability of voting in election s_t . Since eligibility does not mean voting, equation (1) yields the intent-to-treat (ITT) effect of voting in the same type of election as the election in which one was first eligible to vote.

⁵Continuing into tertiary education is also, strictly speaking, post-treatment. We thus employ an additional specification, based on a recoding of education, collapsing in one category all respondents with at least a 12-grade diploma or equivalent. The results appear in Figures C8, C9, C10, and C11 and are similar to those presented in the main text.

⁶We consider age and level of education as necessary controls because they are both strong predictors of turnout, and their distribution in the quasi-experimental groups in our sample pushes the outcomes in the opposite direction of the treatment. For this reason, failing to account for them biases the coefficients toward zero. For example, individuals first eligible for Midterm elections whose outcome is measured after Presidential elections are older than individuals first eligible for Midterm elections whose outcome is measured after Midterm elections. We would expect our treatment (similarity of election type between current and first eligible election) to increase turnout among the latter group. But this group is also younger, which should decrease their probability of voting—since we know that age is a strong predictor of turnout. A similar pattern happens with education. Individuals first eligible for Midterm elections whose outcome is measured after Presidential elections are more educated than individuals first eligible for Midterm elections whose outcome is measured after Midterm elections.

⁷The results from these analyses appear in Figures C6, C7, C12, and C13 of the online Appendix.

Equation (1) does not distinguish between types of elections. This is done in equation (2), which interacts P_s with $Same_{is}$.⁸

$$Vote_{i,s>1} = \sum_{s=2}^{15} \alpha_s E_s + \lambda_1 Same_{is} + \delta P_s + \lambda_2 Same_i P_s + v_i \tag{2}$$

Now, λ_1 yields the effect of same-type election for Midterm elections and $\lambda_1 + \lambda_2$ denote the same effect for Presidential elections. Finally, we decompose the same-type election effects across individuals' life trajectory. Equation (3) does this, by interacting $Same_i$ with each pair of elections in respondent's life span, up until election pair $s = 15$:

$$Vote_i = \sum_{s=2}^{15} \alpha_s E_s + \lambda_1 Same_{is} + \sum_{s=2}^{15} \delta_s Same_i E_s + \delta P_s + X\beta_2 + u_i \tag{3}$$

For each pair of election, $2 \leq s \leq 15$, $\lambda_1 + \delta_s$ yields the effect of same-type-eligibility on the probability of voting in election s . For all elections $s > 15$, the effect is summarized by λ_1 . Equation (4), below, further distinguishes between each type of election:

$$V_i = \sum_{s=2}^{15} \alpha_s E_s + \lambda_1 Same_{is} + \delta P_s + \sum_{s=2}^{15} \beta_s E_s Same_i + \sum_{s=2}^{15} \theta_s E_s P_s + \lambda_2 Same_i P_s + \sum_{s=2}^{15} \zeta_s E_s P_s Same_i + X\beta_2 + u_i \tag{4}$$

Now, for each $2 \leq s \leq 15$, $\lambda_1 + \beta_s$ yields the same-type-of-election effect for all Midterm elections while $\lambda_1 + \beta_s + \lambda_2 + \zeta_s$ yields the same-type-of-election effect on turnout in Presidential elections.

2. Results

The left-hand panel on Figure 3 displays the main findings. Each entry denotes the effect of same-type-eligibility on future turnout. The middle entry comes from equation (1), while the left and right ones come from equation (2). We find a positive effect of almost three percentage point increase in the probability to vote in election s_n , if this election is of same type as election s_1 . Our estimates are of the intent-to-treat, the effect of simply being eligible to vote in one election type rather than another. Since previous work indicates that turnout among first-time-eligible voters is around 60 percent in US Presidential elections (Dinas, 2014), this means that the effect for compliers—i.e., individuals who actually vote in the first election for which they are eligible—would be around five percentage points ($3/0.6 = 5$). Importantly, as shown in the left- and right-most entries of the graph, this effect seems to be of similar magnitude either when the current election is Presidential or Midterm.

These results are not directly comparable with other studies looking at early-election effects (Meredith, 2009; Franklin and Hobolt, 2011; de Kadt, 2017), because such studies compare eligibles with non-eligibles—in other words, their treatment groups differ in terms of number of elections in which they were previously eligible. By contrast, because our main focus is on testing the contextual-continuity hypothesis against the salience hypothesis, we compare individuals in terms of type of first-eligible election. That said, our effect size is comparable to theirs. For example, Meredith (2009) estimates that previous voting experience increases participation by around five percentage points, an effect size that is very similar to ours.

⁸We abuse notation between equations to retain simplicity.

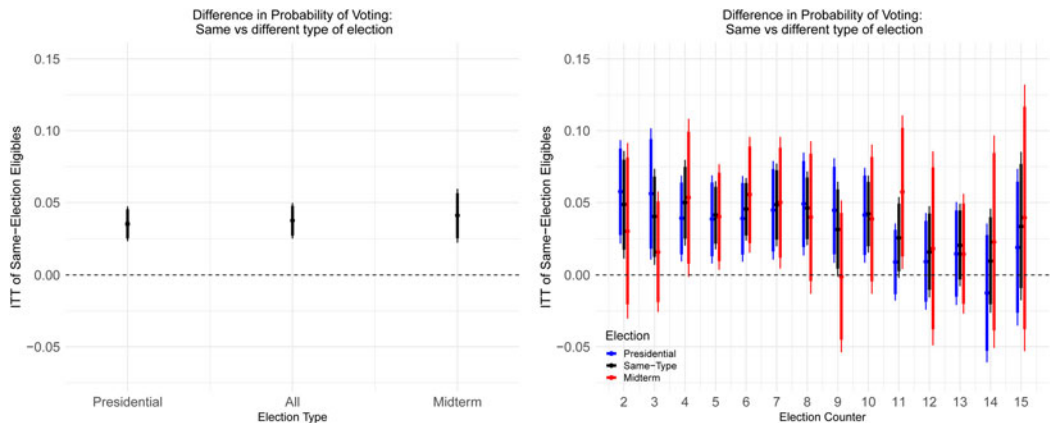


Figure 3. Effect of experiencing the same type of election as the first eligible election on the probability of voting. *Notes:* The vertical spikes denote the 90 percent (thick line) and 95 percent (thin line) confidence intervals. Effects are broken down into the overall effect of same-type eligibility, the effect of being first eligible to vote for a Presidential election on voting in subsequent Presidential elections; and the effect of being first eligible to vote for a Midterm election on voting in subsequent Midterm elections. The analyses control for age, age squared, level of education, female respondents (dummy), and white respondents (dummy).

While providing a summary of the average effect of same-type-eligibility on future turnout, this analysis does not allow us to gauge how much these effects last along the life trajectory. Are these effects short-lived or do they persist across a long span of electoral history? The right-hand panel of Figure 3 tries to answer this question, decomposing the effects across the first 15 pairs of elections encountered in individuals' life span. The black entries denote the overall effect of same-type eligibility, while the blue and red entries decompose this effect into Presidential and Midterm elections respectively. Thus, the estimates denoted in black stem from equation (3), while the blue and red entries stem from equation (4). Three observations are worth making. First, as expected, all estimates carry now higher levels of uncertainty, especially when it comes to Midterm elections, which constitute one-third of the total number of observations. Second and more importantly, despite some fluctuation, the effects seem to be remarkably persistent. Same-type eligibility increases turnout up until the tenth pair of elections individuals encounter. Although the effects decay afterwards, they seem to hold for approximately four decades, ranging between two and five percentage points. Third, the effects are once again symmetric across Presidential and Midterm elections. In other words, same-type-eligibility effects do not seem to be driven only by eligibility in Presidential elections, as the saliency hypothesis would lead us to expect. Rather, and consistent with the contextual continuity hypothesis, the effects are of same direction and comparable magnitude when the same-type election is a Midterm election.

As discussed by Coppock and Green (2016, p. 1058), an alternative interpretation of these results is that instead of denoting the persistent effect of early electoral experiences, they may be driven by differential exposure to party targeting campaigns. If parties' voter files differentiate between types of elections, Midterm eligibles may be contacted more often than Presidential eligibles in Midterm elections, and vice versa. To see whether this is the case, we use as outcome a question available in most of the ANES surveys, asking whether respondents were contacted by the parties during the campaign. Figure C1 in the online Appendix presents the results. We find no evidence that same-type eligibility significantly increases party targeting.

We try to further assess the mechanism driving these effects by examining a complementary implication stemming from the contextual continuity hypothesis. If it is the case that voting habits are linked to the election type that generated them, we should find individuals to be more interested in elections of the type they were first eligible to vote in. We thus replicate the

exercise above using interest in the election as the outcome variable. The results are shown in Figure 4. The pattern is quite similar to the one observed in Figure 3, especially when it comes to the overall effects. Respondents seem to be between two and five percentage points more likely to be interested in an election if this is of the same type as their first eligible election. Moreover, as with the findings shown in Figure 3, the effect seems to be unaffected by whether the first eligible election was a Presidential or a Midterm one. The only substantive difference between the two figures is the degree of persistence in the overall effects. Same-type-eligibility effects appear to be considerably more short-lived when it comes to interest in the election than when it comes to actual voting. This divergence could be explained by the habitual nature of voting: electoral interest boosts turnout in first elections but is not much needed once such habits are formed.

The online Appendix provides robustness checks for these analyses. In the first place, Figures C2 and C3 in the online Appendix provide a placebo test that moves the treatment group by one year, so that the resulting treatment and control groups retain their age differences while now having equal probability of encountering an election of same type as first eligible election. We find no discernible pattern, either when looking at turnout (Figure C2) or when looking at interest in the election (Figure C3). Figures C4 and C5 replicate the main analyses shown in Figures 3 and 4 respectively, this time adding state-fixed effects. The results remain substantively unchanged. Figures C6 and C7 replicate the analyses without additional control variables, while C12 and C13 extend the list of covariates to contemporaneous demographics. Both sets of results point to the same conclusions as those presented here. Figures C10 and C11 extend the analysis with no additional controls by further using a coding scheme for education that does not distinguish between high-school and above, to permit variation in education that has manifested itself only by the time the first eligible election takes place. The results are similar to those presented above. Moreover, since the election takes place in November, we treat all individuals coming of age in an election year as eligible to vote in that election. Figures C14 and C15 relax this assumption, presenting the results when excluding everyone becoming of age in an election year. The results remain substantively very similar to those presented in the main analysis. Finally, Figures C16, C17, and C18 replicate the main findings displayed in the left panel of Figure 3, after excluding one election year at a time. As expected, the point estimates are now

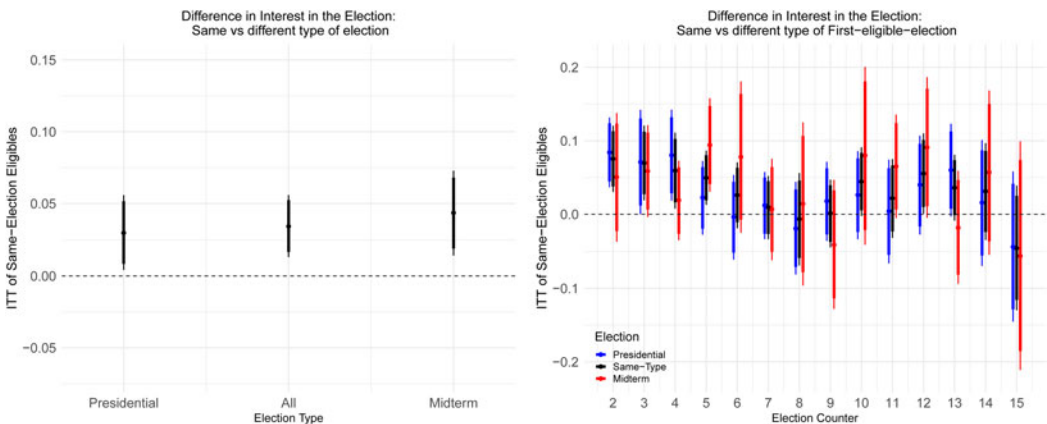


Figure 4. Effect of experiencing the same type of election as the first eligible election on interest in the election. Notes: The vertical spikes denote the 90 percent (thick line) and 95 percent (thin line) confidence intervals. Effects are broken down into the overall effect of same-type eligibility, the effect of being first eligible to vote for a Presidential election on interest in subsequent Presidential elections; and the effect of being first eligible to vote for a Midterm election on interest in subsequent Midterm elections. The analyses control for age, age squared, level of education, female respondents (dummy), and white respondents (dummy).

less precisely estimated, but remain substantively very similar. Figures C19, C20, and C21 repeat the same exercise to replicate the findings for interest in elections (Figure 4). Again, the results remain very similar after excluding one survey at a time.

3. Policy implications

If the type of first eligible election matters, it should also help us understand the effects of policies designed to reduce voting age. Reducing the age threshold of voting eligibility has become increasingly popular as a way of increasing youth political knowledge (Rosenqvist, 2020) and locking young individuals into voting habits (Franklin, 2020). If such policies change the type of first-eligible elections, they may also have hitherto neglected side-effects on the probability of voting for subsequent elections of different types. To shed light on these effects, we leverage the laws of 1972 reducing the voting age from 21 to 18 years old that affected most US States.

The 1972 law made voters born in years such as 1963, 1967, or 1971 first eligible to vote in midterm elections at age 19. In the absence of the reform, they would have been first eligible to vote in a Presidential election at age 21. On the other hand, people born a year before remain first eligible to vote in Midterm elections regardless of the reform. Comparing people aged 19 and 20 at the time of congressional elections before and after the reform, we can therefore test whether making one group to experience congressional elections first affects future turnout. Doing so also presents the advantage of accounting for potential cohort-specific fixed effects, which was impossible with the previous research design.

To test the effect of switching from a presidential to congressional election as a first eligible voting experience, we use the same data as in the previous analyses and estimate the following equation:

$$\begin{aligned} Voted_{i, s>1} = & \gamma_0 Change_i + \gamma_1 Presidential_i Change_i + \delta Presidential_i \\ & + \left[\sum_{s=2}^{15} \alpha_{0s} E_s \right] + X_i \beta_0 + Presidential_i \left(\left[\sum_{s=2}^{15} \alpha_{1s} E_s \right] + X_i \beta_1 \right) + v_{is} \end{aligned} \quad (5)$$

In the equation, *Change* equals one if the individual is 19 years old in the year of congressional elections and is born after 1955—thereby experiencing a switch in their first eligible election from Presidential to Midterm. Alternatively, *Change* equals zero if the voter turns 19 in the year of congressional elections and is born before 1955 or turns 20 in the year of such elections and therefore always remains first eligible for presidential elections. We also include our *Presidential* dummy to control for turnout differences between these two elections as well as an interaction between *Presidential* and *Change*. This means that γ_0 represents the effect of switching to midterm elections on subsequent midterm election turnout while $\gamma_0 + \gamma_1$ represents the effect of switching to midterm on subsequent presidential elections' turnout. Finally, we include the same covariates and election-counter fixed effects as before in a fully interacted model to allow for these to depend on the type of election voters experience. To verify the robustness of our results, we also run the same regressions with state code fixed effects and their corresponding interactions with the presidential dummy. Results do not change, albeit becoming less precise.

Table 3 shows the findings.⁹ Regardless of model specification, switching one's first election from Presidential to Midterm increases the likelihood of voting in subsequent Midterm elections

⁹In these analyses, we remove women born 21 years before female voting rights were adopted and afro-american born more than 21 years before 1965. We also remove the states that reduced voting rights before 1972: Kentucky, Georgia, Alaska and Hawaii.

Table 3. Effect of changing one’s first election type from Presidential to Midterm on voting for subsequent elections.

	(1)	(2)	(3)
Change	0.0444** (0.0183)	0.0387** (0.0186)	0.0305 (0.0187)
Change x midterm	-0.0935*** (0.0249)	-0.0794*** (0.0263)	-0.0681** (0.0264)
Observations	21,285	21,278	21,278
Necessary controls	Yes	Yes	Yes
Necessary controls interacted	Yes	Yes	Yes
Additional controls	No	Yes	Yes
Additional controls Interacted	No	Yes	Yes
State dummies	No	No	Yes
State dummies interacted	No	No	Yes
Election pair fixed effects	Yes	Yes	Yes
Election pair interacted	Yes	Yes	Yes

Robust standard errors in parentheses
 ***p < 0.01, **p < 0.05, *p < 0.1

by three to four percentage points and decreases the probability of voting in subsequent Presidential elections by between four and five percentage points.

Table C1 in the online Appendix repeats the same exercise focusing on the group who experience a change from Midterm to Presidential elections after the reform. As expected the coefficients are of opposite sign to the ones shown in Table 3.

To make sure that we are not rather measuring the effect of one group now being eligible to vote in an election a year earlier than the other, Table C2 in the online Appendix compares 18 and 19 year olds after the reform who are eligible for the same first election assigning treatment to 19 year olds (models 1 and 2). We find no impact both for congressional and presidential elections. We repeat the exercise comparing 21 and 22 years who became first eligible for the same election pre-1972 (models 3 and 4) and, again, find no effect.

4. Discussion

Drawing upon literature on habit formation, we have challenged the view that early elections exert more powerful effects on future turnout when they are of high salience. Given the contextual continuity between elections that share a type, we argued that one should expect early elections to boost turnout in future elections of the same type. Our analyses show that, indeed, individuals first eligible to vote for a Presidential election are more likely to vote at subsequent Presidential elections; while individuals first eligible to vote for a Midterm election are more likely to vote at subsequent Midterm elections. These results have also non-negligible policy implications: policies that change the first election individuals are eligible to have a similar effect. This finding supports the suspicion put forward by Coppock and Green (2016, p. 1060) that downstream effects of voting experiences are more likely to happen in similar elections.

Our results underline the importance of the electoral context and question the ability of habits formed in previous elections to travel across elections of different types. The conventional view in the literature is that voting in previous elections—if these are sufficiently salient—can induce voting habits that travel to different contexts. Our results paint a more nuanced picture. They suggest that, while participatory habits are formed, such habits are rather constrained. Individuals seem to form election-specific habits that do not necessarily affect elections of different types.

Such conclusion qualifies the view that second-order elections can be damaging for electoral participation—an argument that has fueled criticism of the existence of low salience elections, such as elections for the European Parliament. As our findings suggest, early voting experiences can boost turnout in same-type elections. While saliency is indeed more important in predicting

turnout across elections—hence the overall gap between Midterm and Presidential election turnout—the long-term effects of early voting experiences seem to be affected more by the type than the salience of the first-eligible election.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/psrm.2023.2>. To obtain replication material for this article, please visit <https://doi.org/10.7910/DVN/YGPAW>.

References

- Aldrich JH, Montgomery JM and Wood W** (2011) Turnout as a habit. *Political Behavior* **33**, 535–563.
- Coppock A and Green DP** (2016) Is voting habit forming? New evidence from experiments and regression discontinuities. *American Journal of Political Science* **60**, 1044–1062.
- de Kadt D** (2017) Voting then, voting now: the long-term consequences of participation in South Africa’s first democratic election. *The Journal of Politics* **79**, 670–687.
- Denny K and Doyle O** (2009) Does voting history matter? Analysing persistence in turnout. *American Journal of Political Science* **53**, 17–35.
- Dinas E** (2012) The formation of voting habits. *Journal of Elections, Public Opinion and Parties* **22**, 431–456.
- Dinas E** (2014) Does choice bring loyalty? Electoral participation and the development of party identification. *American Journal of Political Science* **58**, 449–465.
- Dinas E and Riera P** (2018) Do European parliament elections impact national party system fragmentation?. *Comparative Political Studies* **51**, 447–476.
- Franklin MN** (2004) *Voter turnout and the dynamics of electoral competition in established democracies since 1945*. Cambridge MA: Cambridge University Press.
- Franklin MN** (2020) *Consequences of Lowering the Voting Age to 16: Lessons from Comparative Research*. Cham: Springer International Publishing, pp. 13–41.
- Franklin MN and Hobolt SB** (2011) The legacy of lethargy: how elections to the European parliament depress turnout. *Electoral Studies* **30**, 67–76.
- Gerber AS, Green DP and Shachar R** (2003) Voting may be habit-forming: evidence from a randomized field experiment. *American Journal of Political Science* **47**, 540–550.
- Green DP and Shachar R** (2000) Habit formation and political behaviour: evidence of consuetude in voter turnout. *British Journal of Political Science* **30**, 561–573.
- Hall ET** (1969) *Ecological Psychology: Concepts and Methods for Studying the Environment of Human Behavior*. Standord: Stanford University Press.
- Ji MF and Wood W** (2007) Purchase and consumption habits: not necessarily what you intend. *Journal of Consumer Psychology* **17**, 261–276.
- Larcom S, Rauch F and Willems T** (2017) The benefits of forced experimentation: striking evidence from the London underground network. *The Quarterly Journal of Economics* **132**, 2019–2055.
- Meredith M** (2009) Persistence in political participation. *Quarterly Journal of Political Science* **4**, 187–209.
- Plutzer E** (2002) Becoming a habitual voter: inertia, resources, and growth in young adulthood. *American Political Science Review* **96**, 41–56.
- Rosenqvist O** (2020) Rising to the occasion? Youth political knowledge and the voting age. *British Journal of Political Science* **50**, 781–792.
- Schulte-Cloos J** (2019) The effect of European parliament elections on political socialisation. *Journal of European Public Policy* **26**, 1094–1111.
- Wood W and Neal DT** (2007) A new look at habits and the habit-goal interface. *Psychological Review* **114**, 843.
- Wood W, Tam L and Witt MG** (2005) Changing circumstances, disrupting habits. *Journal of Personality and Social Psychology* **88**, 918.