

Differential Registration Bias in Voter File Data: A Sensitivity Analysis Approach*

Brendan Nyhan[†]

Christopher Skovron[‡]

Rocío Titiunik[§]

January 12, 2016

Abstract

The widespread availability of voter files has improved the study of political participation in American politics, but the fundamental unavailability of voting-eligible population (VEP) counts remains an obstacle to calculating turnout rates. Many observational studies simply calculate turnout among registrants, which may lead to invalid inferences if registration rates differ between treatment and control groups. We introduce a sensitivity analysis to assess the potential bias induced by differential registration. Our approach is most helpful when only a cross-section of registrants is available, but is also useful when only the subpopulation registered before treatment is included or when a proxy VEP measure is used as denominator. We illustrate our approach with two studies of how voting eligibility affects subsequent turnout among young voters. In both cases, eligibility appears to decrease turnout among registrants, but our approach shows that these negative effects are highly sensitive to differential registration bias.

*We thank Matias Cattaneo, Alexander Coppock, Simon Chauchard, Kevin Collins, Jeff Friedman, Brian Greenhill, Michael Herron, Yusaku Horiuchi, Jeremy Horowitz, Dean Lacy, Jacob Montgomery, Phil Paolino, Jason Reifler, Daniel Smith, Brad Spahn, Thomas Zeitzoff and seminar participants at the UCLA American Politics Workshop for helpful feedback; Bob Blaemire at Catalist for his assistance in procuring our data; and John Holbein and Sunshine Hillygus for providing replication data. All errors are our own. Nyhan acknowledges support from the Robert Wood Johnson Foundation's Scholars in Health Research Program. Skovron acknowledges funding from the NSF Graduate Research Fellowship Program. Titiunik acknowledges financial support from the National Science Foundation (SES 1357561).

[†]Assistant Professor, Department of Government, Dartmouth College, nyhan@dartmouth.edu

[‡]Ph.D. Candidate, Department of Political Science, University of Michigan, cskovron@umich.edu.

[§]Assistant Professor, Department of Political Science, University of Michigan, titiunik@umich.edu.

The widespread availability of digital voter files has changed the study of voting behavior in American politics. These files offer data on a vast population of U.S. citizens while avoiding the social desirability bias and low statistical power that plague survey studies of self-reported turnout, enabling new studies of how factors as disparate as majority-minority districts (Barreto, Segura and Woods 2004), minority candidates (Barreto 2007; Fraga 2015), genetic similarity (Fowler, Baker and Dawes 2008), and early voting registration (Holbein and Hillygus 2015) affect turnout at the individual or aggregate level. These data have been found to be of high quality, especially when cleaned and aggregated by third party firms like Catalist (Ansolabehere and Hersh 2012; Hersh 2015).

Scholars of political behavior use registered voter lists in at least two different ways. In what we call *experimental studies of registrants*, registration records are used to identify a subset of citizens among which a treatment is randomly assigned. For example, studies of get-out-the-vote (GOTV) campaigns typically start with a list of registered voters, randomly assign each citizen in the list to be encouraged to vote (or not), and use future registration files to measure subsequent turnout (e.g., Gerber, Green and Larimer 2008; Citrin, Green and Levy 2014). In this case, the registered population is the universe from which a sample is selected to conduct an experiment; registration occurs *before* the treatment is assigned.

In contrast, in what we call *non-experimental turnout studies*, researchers are typically interested in the effect of a non-experimental treatment on voter turnout (or partisan registration) and use voter registration files as the source of outcome data. In this case, the treatment of interest is likely to affect both the likelihood of registration and the likelihood of turning out to vote (or the likelihood of choosing to register as a partisan). For example, the presence of a minority candidate on the ballot may lead to higher minority registration as well as higher minority turnout.

These two cases differ crucially in their study populations. In experimental studies of registrants, the population of interest is the initial group of registrants, which is collected before the treatment is assigned. In contrast, non-experimental turnout studies consider treatments that could affect both registration and voter turnout decisions. In these studies, the decision to register may be a consequence of the treatment itself, creating the potential for a form of post-treatment bias (Rosenbaum 1984) sometimes known as “endogenous selection bias” (Elwert and Winship 2014).

If a treatment affects registration rates as well as voter turnout (the outcome of interest), a simple comparison of turnout rates between the treated and control groups among registered voters could lead to mistaken inferences.

We formally show the threat that conditioning on the population of registrants poses in studies where the treatment of interest affects the likelihood of registration, characterizing the bias that stems from differential registration in treated and control groups. We also develop a novel sensitivity analysis that makes it possible for scholars to assess the robustness of their treatment effect estimates to (unobserved) potential differences in registration rates between treatment and control groups. In addition, our sensitivity method offers a useful complementary approach for all non-experimental turnout (or partisanship) studies where the treatment may itself affect the composition of the registered population.

Our survey of the literature on non-experimental turnout studies indicates that researchers who wish to investigate the turnout or partisanship effects of an intervention that also affects registration tend to adopt one of three strategies. Some studies, which are typically cross-sectional, simply use the registration file as the universe of analysis and calculate turnout or partisanship rates as the proportion of voters or partisans among registrants (e.g., Barreto, Segura and Woods 2004; Barreto 2007; Fraga and Merseth 2016; Hersh 2013; Hersh and Nall 2016). Other studies identify the period when the intervention or treatment is introduced, collect registration files from a period before treatment, and look at the effect of the treatment only for the subpopulation of registrants identified before the treatment (e.g., Barber and Imai 2014; Enos 2016; Fraga 2016). For example, Enos (2016) studies the effect of the demolition of housing projects on white turnout only among individuals who were registered before and after the demolition. Finally, other scholars rely on voter files to calculate the numerator of interest (number of voters or partisans), but calculate turnout or partisanship rates using a measure of the potential electorate obtained from a secondary source as the denominator (Meredith 2009; Fraga 2015; Cepaluni and Hidalgo 2016). For example, Fraga (2015) uses the American Community Survey's estimates of the citizen voting-age population to approximate the population of eligible voters.

Our sensitivity approach is useful for each of these types of studies. In the first case, when the only data available is a cross-section of registered citizens, treatment effects on turnout or parti-

sanship can change dramatically if there is differential registration between treatment and control groups. Our approach is most useful in this case, offering a concrete measure of the robustness of treatment effects that are estimated on registrants only.

In the second case, when only a subset of the registered population is studied before and after the intervention, scholars must exclude from the analysis all citizens who register after the intervention—a potentially very large group. For example, the before-and-after analysis in Barber and Imai (2014) includes only 25% to 29% of registered voters (see p. 26-27). Moreover, this approach is valid only if the likelihood of appearing in the registration file after the intervention is not affected by the intervention itself. Our sensitivity approach allows researchers to study the robustness of treatment effects for the entire population, complementing analyses that focus only on the subset of the population that appears before and after.

Finally, our approach is also useful in the third case when a secondary source is used as the denominator of turnout (or partisanship) rates. It is well known that estimated turnout rates are sensitive to the specification of the total eligible population (McDonald and Popkin 2001). In the U.S., census counts of the voting-eligible population (VEP) are not available. Common choices to approximate the VEP are the voting-age population (VAP) or the citizen voting-age population (CVAP), but these measures are imperfect approximations: both include ineligible populations such as disenfranchised felons; the VAP includes noncitizens; and the CVAP is estimated based on surveys—not census counts—and is unavailable for the smallest census geographies. Thus, in studies that calculate the denominator from secondary sources, our approach is useful to determine whether the differential registration that would change the conclusions of the study are within the plausible margin of error of the VEP estimate.

We illustrate the challenge of differential registration and our sensitivity approach with two studies of political socialization, a research area that has recently begun to use quasi-experimental techniques to estimate the effect of initial election eligibility on subsequent voter turnout and other behaviors (e.g., Meredith 2009; Mullainathan and Washington 2009; Dinas 2012, 2014; Coppock and Green 2015; Holbein and Hillygus 2015). These studies typically compare voter turnout and political attitudes among individuals whose 18th birthday fell close to a previous general election. We use a similar research design that compares turnout rates between young citizens who were

narrowly eligible or ineligible to vote in a prior election.

Our first study is an original analysis of a sample of registrants from 42 U.S. states and the District of Columbia who were born within four days of the election eligibility cutoff—a far narrower window than previous studies. Our findings illustrate the need to consider the potential bias induced by conditioning on registration. When we naïvely use registered voters in the voter file as the denominator for our turnout analysis, we find results that largely contradict previous studies—initial election eligibility appears to sometimes *reduce* subsequent turnout or have no effect. However, we show that this result is highly sensitive to differing registration rates between the two groups. Consistent with this sensitivity analysis, when we instead use birth totals as the denominator, we find that initial eligibility *increases* subsequent turnout. This application shows that differential registration between treatment and control groups can severely bias non-experimental turnout studies when turnout rates are calculated as a proportion of total registration.

Our second study is a re-analysis of the Florida analysis in Holbein and Hillygus (2015), where we again examine the effect of voting eligibility on subsequent turnout. Like the first study, eligibility is found to negatively affect turnout in later elections among young voters when turnout rates are calculated among registered voters. However, a sensitivity analysis of the data indicates that even a slightly higher registration rate in the treatment group would reverse these results. When we use birth counts to approximate the total population, we find that the registration rate is higher in the control group than in the treatment group in the window around the eligibility cutoff originally used by Holbein and Hillygus, but this finding is reversed when a larger window is considered. Our analysis suggests that the negative effects reported by Holbein and Hillygus are thus highly sensitive to differential registration and consistent with true null or even positive effects. This case illustrates the usefulness of our sensitivity method in cases where approximations of the voting eligible population are imperfect.

These results suggest that scholars who study turnout based on voter files should complement their analysis with a rigorous sensitivity analysis; even small differences in registration rates between treatment and control groups can reverse the conclusions of observational turnout studies. Our goal is not to claim that past or current studies based on voter files are incorrect or misleading, but instead to help raise awareness of this difficult inferential issue and provide researchers with a

simple approach to assess the robustness of their findings.

Defining and assessing differential registration bias

We begin by illustrating the problem of differential registration bias in Table 1 with a hypothetical example of an election in which 10,000 citizens are in the treatment group and 10,000 citizens are in the control group. In addition, the treatment is associated with higher rates of registration—5,000 individuals in the treatment group register compared with 4,000 in the control group (column 2)—but has no effect on turnout (columns 3 and 5). If we simply compare turnout rates among registered voters using voter file data (column 4), we would conclude that turnout rates are *lower* among treated voters ($2500/5000=50\%$ versus $2500/4000=62.5\%$ among controls). However, the true turnout rate of 25% is identical in both groups (column 5).

Table 1: Hypothetical example of differential registration bias

	Total (1)	Registered (2)	Voted (3)	Turnout (% reg.) (4)	Turnout (% pop.) (5)
Treated	10,000	5,000	2,500	50%	25%
Control	10,000	4,000	2,500	62.5%	25%

This problem would of course be avoided if we could calculate turnout rates using the total voting-eligible population in each group rather than the total number of *registered voters*. When possible, the simplest solution to the problem of differential registration is to obtain the missing population totals from alternative data sources. However, as we discussed above, the necessary data is typically either unavailable or imperfectly approximated.

A formal sensitivity analysis approach

We formalize a sensitivity analysis approach that can be implemented in cases where the needed quantities—the total population in the treatment and control groups—are unknown or imprecisely estimated. Specifically, our approach identifies the differential in registration rates between the treatment and control groups that would reduce the observed difference in turnout rates among registered voters to zero. This method provides a measure of the vulnerability of a treatment effect

estimate to differential registration bias when population totals are missing. It can also be used to corroborate findings when population totals from alternative data sources are likely to be estimated with error, which is a pervasive phenomenon because precise counts of the voting-eligible population are not directly available (McDonald and Popkin 2001). Moreover, as we discuss in the conclusion, our approach could also be extended to other types of data that share the characteristics of turnout data (unknown population totals and values of missing data known with certainty).

In the example illustrated in Table 1 above, we assumed that the true population counts were known. We now assume that total population counts for the treatment and control groups are not available and that researchers are working with a voter file that includes only registered voters.¹ This scenario is illustrated in Table 2, where for each group we have the total registration (column 2) and the total number of voters (column 3), which allows us to calculate turnout rates as a proportion of total registration in each group (column 4). However, letting the subscripts T and C denote the treatment and control groups, respectively, the total population counts in each group, which we denote by P_T and P_C , are not available. As a consequence, the true turnout rates—the ratio of voters to the total population in each group—are also unavailable, creating the risk of a mistaken inference if differential registration bias is present.

Table 2: Hypothetical illustration of sensitivity analysis

	Overall population				
	Total	Registered	Voted	Turnout (% reg.)	Turnout (% pop.)
	(1)	(2)	(3)	(4)	(5)
Treated	P_T	5,000	2,500	50%	$2,500/P_T$
Control	P_C	4,000	2,500	62.5%	$2,500/P_C$

We introduce some additional notation. We let R_T and R_C denote the total registration counts in the treatment and control groups, respectively, and V_T and V_C denote the total numbers of voters who turned out to vote in each group. The desired—but unavailable—turnout rates are thus

$$T_T^{\text{Pop}} = \frac{V_T}{P_T} \quad \text{and} \quad T_C^{\text{Pop}} = \frac{V_C}{P_C},$$

where the superscript Pop denotes that these turnout rates are calculated as a proportion of the total

¹For simplicity, we assume the entire registration file is available. However, our argument applies directly to the case where only a random sample from the voter file is available (as in Study 1 below).

population. Henceforth, we refer to T_T^{Pop} and T_C^{Pop} as the *turnout-to-population* rates or simply the *true turnout* rates. Without observing the total population counts P_T and P_C , we cannot calculate these turnout-to-population rates directly.

We define the registration rates $r_T = R_T/P_T$ and $r_C = R_C/P_C$ in the treatment and control groups, respectively. Given these rates, the total registration counts, R_T and R_C , can be expressed as $R_T = r_T \times P_T$ and $R_C = r_C \times P_C$, and we can express the unknown total population as total registration divided by the probability of registration:

$$P_T = R_T/r_T \quad \text{and} \quad P_C = R_C/r_C.$$

We then define the *turnout-to-registration* rate in each group, T_T^{Reg} and T_C^{Reg} , as follows:

$$T_T^{\text{Reg}} = \frac{V_T}{R_T} \quad \text{and} \quad T_C^{\text{Reg}} = \frac{V_C}{R_C}.$$

Given these definitions, we can express the desired turnout-to-population rates as

$$T_T^{\text{Pop}} = \frac{V_T}{P_T} = \frac{V_T}{R_T/r_T} = T_T^{\text{Reg}} \times r_T \quad \text{and} \quad T_C^{\text{Pop}} = \frac{V_C}{P_C} = \frac{V_C}{R_C/r_C} = T_C^{\text{Reg}} \times r_C. \quad (1)$$

In words, the true turnout rate is the turnout-to-registration rate adjusted (multiplied) by the registration rate in each group. In applications where researchers have access to the voter file but not the total population, the registration rates r_T and r_C are unknown. Our sensitivity analysis considers how large the difference between r_T and r_C would have to be to generate the observed difference in turnout-to-registration rates when the true difference in turnout-to-population rates is zero.

Imagine that $T_T^{\text{Reg}} < T_C^{\text{Reg}}$, which means that the turnout-to-registration rate is smaller in the treatment group than in the control group (as in the two applications we consider below). Because the registration rates for the two groups r_T and r_C are unknown, $T_T^{\text{Reg}} < T_C^{\text{Reg}}$ is not enough to conclude that $T_T^{\text{Pop}} < T_C^{\text{Pop}}$.² In other words, a lower (higher) turnout-to-registration rate in the treatment group does not necessarily imply that this group has a lower (higher) turnout-to-population rate. However, given the observed difference $T_T^{\text{Reg}} - T_C^{\text{Reg}}$, we can estimate how different the reg-

²Likewise, we can't conclude that $T_T^{\text{Pop}} > T_C^{\text{Pop}}$ or that $T_C^{\text{Pop}} = T_T^{\text{Pop}}$ in this case either without knowing r_T and r_C .

istration rates would have to be between the treatment and the control groups for this difference to be observed when there is no difference in turnout-to-population rates (i.e., $T_T^{\text{Pop}} - T_C^{\text{Pop}} = 0$).

To illustrate our approach, we define Δ as the treatment-control difference in true turnout rates:

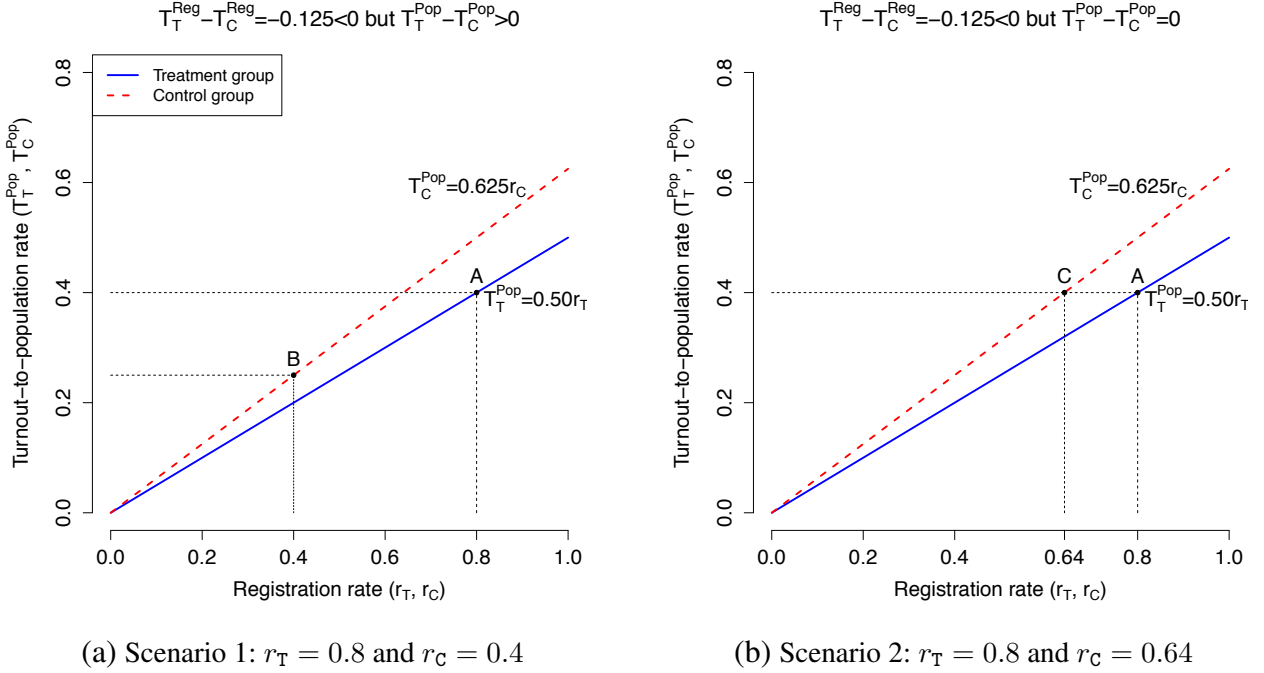
$$\Delta = T_T^{\text{Pop}} - T_C^{\text{Pop}} = (T_T^{\text{Reg}} \times r_T) - (T_C^{\text{Reg}} \times r_C).$$

We note two points. First, if $r_T = r_C = 1$, the expression simplifies to $\Delta = T_T^{\text{Reg}} - T_C^{\text{Reg}}$. In other words, if everyone in the treatment and control groups registers, the ratio of voters to registrants is of course identical to the ratio of voters to total population and there are no complications. Second, if turnout rates are identical between groups but not everyone votes ($r_T = r_C = r \neq 1$), the unknown turnout share difference simplifies to $\Delta = r(T_T^{\text{Reg}} - T_C^{\text{Reg}})$. Since $0 \leq r \leq 1$, the sign of Δ will be equal in this case to the sign of $T_T^{\text{Reg}} - T_C^{\text{Reg}}$ and $|\Delta| \leq |T_T^{\text{Reg}} - T_C^{\text{Reg}}|$; moreover, it is straightforward to explore how Δ changes as r varies from 0 to 1.

We are interested in the more general case in which both r_T and r_C are nonzero and $r_T \neq r_C$. Figure 1 illustrates how a negative “treatment effect” on turnout-to-registration rates can be observed even when the true difference in the turnout-to-population rates is null or even positive. Using the functions defined in Equation (1), the figure plots the turnout-to-population rates (T_T^{Pop} , T_C^{Pop}) against the registration rates (r_T , r_C) separately for each group holding turnout-to-registration rates (T_T^{Reg} , T_C^{Reg}) fixed. We adopt the scenario in Tables 1 and 2 where $T_T^{\text{Reg}} = 0.50$ and $T_C^{\text{Reg}} = 0.625$. Thus, we plot the linear functions $T_T^{\text{Pop}} = 0.50 \times r_T$ and $T_C^{\text{Pop}} = 0.625 \times r_C$. Since the slope in the control group (0.625) is higher than the slope in the treatment group (0.50), the dotted line (control) is always above the solid line (treatment). In other words, Figure 1 fixes $T_C^{\text{Reg}} = 0.50$ and $T_C^{\text{Reg}} = 0.625$, which means that the turnout-to-registration difference is always negative (-0.125).

Imagine first that the probability of registration in the treatment group is 0.8 ($r_T = 0.8$) and the probability of registration in the control group is 0.4 ($r_C = 0.4$). In this case, the true turnout rate in the treatment group (T_T^{Pop}) of 0.4 is obtained from the y-coordinate of point A on the solid line in Figure 1(a), and the true turnout rate of 0.25 in the control group (T_C^{Pop}) is obtained from the y-coordinate of point B on the dashed line. Under this scenario, $\Delta = T_T^{\text{Pop}} - T_C^{\text{Pop}} = 0.4 - 0.25 = 0.15$. In other words, the treatment effect on the true turnout-to-population rate is positive despite the turnout-to-registration rate being lower in the treatment than in the control group. Figure 1(b)

Figure 1: True turnout rates as a function of registration rate



Note: We assume $T_T^{Reg} = 0.50$ and $T_C^{Reg} = 0.625$ in both scenarios.

shows a different scenario in which we assume that $r_T = 0.8$ and $r_C = 0.64$. In this case, both T_T^{Pop} and T_C^{Pop} are equal to 0.4 and thus $\Delta = 0.4 - 0.4 = 0$, though the difference in turnout-to-registration rates is still negative (again, treatment slope is 0.5 and control slope is 0.625).

We propose a sensitivity analysis to explore how our inferences are affected by differential registration bias under these circumstances (the most likely scenario in practice). First, we define the *differential registration factor* k as the ratio of the registration rate in the treatment group to the registration rate in the control group:

$$k = \frac{r_T}{r_C}$$

We have assumed that both r_T and r_C are nonzero so this ratio is well-defined. Moreover, because r_T and r_C are rates or probabilities, they are both less than (or at most equal to) 1, which means that $k \in (0, \infty)$. Furthermore, since $r_T/k = r_C$ and r_C is a rate, k must satisfy the restriction $0 < r_T/k \leq 1$; that is, the smallest value that k can take is r_T .³ Our sensitivity analysis explores, for a given treatment group registration rate r_T , how large the differential registration factor k can

³In other words, values of $k \in [0, r_T)$ are not allowed because they would imply $r_C \geq 1$.

be before the implied difference in turnout-to-population rates is zero or has the opposite sign from $T_T^{\text{Reg}} - T_C^{\text{Reg}}$, the observed difference in turnout-to-registration rates.

Given our definition of k , we can express Δ , the treatment-control difference in true turnout rates, as a function of the treatment group registration rate r_T and the ratio of treatment/control registration rates k :

$$\Delta(r_T, k) = (T_T^{\text{Reg}} \times r_T) - (T_C^{\text{Reg}} \times \frac{r_T}{k}) = r_T(T_T^{\text{Reg}} - \frac{T_C^{\text{Reg}}}{k})$$

where we make the arguments k and r_T explicit. Thus, for any nonzero value of r_T , we can calculate the value of k under which a zero difference in true turnout rates between the treatment and control groups would result in the observed difference in turnout-to-registration rates. Since we assume $r_T > 0$, we find this value, which we call k^* , as the solution to $(T_T^{\text{Reg}} - \frac{T_C^{\text{Reg}}}{k^*}) = 0$, leading to

$$k^* = \frac{T_C^{\text{Reg}}}{T_T^{\text{Reg}}}.$$

By definition, $\Delta(r_T, k^*) = 0$ for any $0 < r_T \leq 1$.

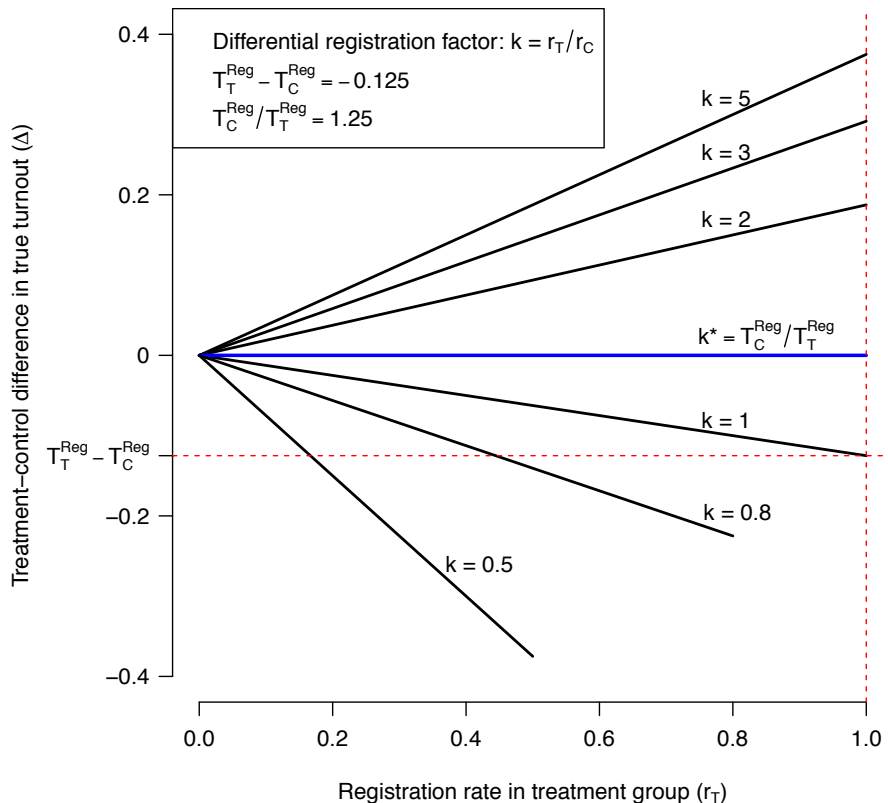
This sensitivity analysis approach thus allows us to explore how the true turnout difference Δ varies with observed turnout-to-registration rates under different assumptions about registration rates for the treatment and control groups (r_T and r_C). In particular, we can calculate k^* , the pattern of differential registration that would be required to produce the observed difference in turnout-to-registration rates if there is no difference in turnout-to-population rates.

We again illustrate the procedure with the hypothetical example presented in Tables 1 and 2. In that example, $T_T^{\text{Reg}} = 2500/5000 = 0.5$ and $T_C^{\text{Reg}} = 2500/4000 = 0.625$, so the difference in turnout rates among registered voters is negative ($T_T^{\text{Reg}} - T_C^{\text{Reg}} = -0.125$). In this case, $k^* = 0.625/0.50 = 1.25$, which means that the observed difference in turnout-to-registration rates could occur under a zero (or positive) difference in true turnout rates if the probability of registration were (more than) 25% higher in the treatment than in the control group.

Figure 2 visualizes how the true turnout difference can vary with k —the ratio of the treatment/control registration rates—for a given difference in turnout-to-registration rates. We plot $\Delta(r_T, k)$ as a function of r_T for differing values of k (which implicitly fixes r_C). The y-axis is the

difference in turnout-to-population rates between the treatment and control groups and the x-axis is the registration rate in the treatment group. As illustrated by the blue curve, when $k = k^* = 1.25$, the difference in true turnout rates is zero for every value of r_T . Moreover, when the registration rate is equal in both groups and thus $k = 1$, the observed difference in turnout-to-registration rates ($T_T^{\text{Reg}} - T_C^{\text{Reg}} = -0.125$) is equal to the true difference in turnout rates if $r_T = 1$.

Figure 2: Difference in true turnout rates as function of k and r_T



Note: We assume $T_T^{\text{Reg}} = 0.50$ and $T_C^{\text{Reg}} = 0.625$.

We can explore sensitivity further in our hypothetical example by calculating the turnout rate difference Δ for values of the registration ratio k above and below the threshold k^* . For $k > k^* = 1.25$, the difference in true turnout is positive and the sign of the turnout-to-registration rate difference is reversed. For $1 \leq k < k^* = 1.25$, the true turnout effect is negative but smaller in absolute value than the difference in turnout-to-registration rates for all r_T . Finally, for $k < 1$, the true turnout effect is negative and can be larger in absolute value than the difference in turnout-to-registration rates for a high enough r_T .

Incorporating prior knowledge of differential registration

The above approach allows researchers to assess a worst-case scenario in which differential registration bias produces the observed difference in turnout-to-registration rates when the true turnout-to-population effect is zero. However, in some applications, the value taken by k^* may not be plausible or informative. We therefore describe a variant of this approach in which scholars can incorporate prior knowledge about plausible variation in registration rates in assessing the conditions under which their results will hold.

Although the rates of registration r_T and r_C are unknown, scholars often have prior information about the range of plausible values they can take. Imagine that we use a survey estimate of the registration rate in the overall U.S. population as a guess for r_T , the true registration rate in the treatment group. We call this guess \tilde{r}_T . Our concern is that r_T differs from r_C —i.e., that the treatment has an effect on registration. In most cases, researchers will be able to offer some prior knowledge about how large the differential registration effect is likely to be and rule out extreme values of $k = r_T/r_C$. Imagine that our guess for the differential registration factor is \tilde{k} . We can use \tilde{r}_T and \tilde{k} to calculate the guessed treatment-control difference in true turnout rates, $\tilde{\Delta} = \tilde{r}_T(T_T^{\text{Reg}} - T_C^{\text{Reg}}/\tilde{k})$. If $\tilde{\Delta}$ is of the same sign as $T_T^{\text{Reg}} - T_C^{\text{Reg}}$, we can conclude that the observed difference in turnout-to-registration rates is robust to a plausible scenario of differential registration rates based on prior knowledge (which may be less stringent than the worst-case scenario represented by k^*).

Consider our example above in which $T_T^{\text{Reg}} = 0.50$ and $T_C^{\text{Reg}} = 0.625$. Let us assume that our guess for \tilde{r}_T is 0.59, the rate of registration in the overall U.S. population estimated by the U.S. Census Bureau’s Current Population Survey in November 2014. Imagine that, based on prior knowledge, we believe that the treatment of interest is unlikely to increase the registration rate in the treatment group by more than ten percentage points relative to the control group. Since $\tilde{r}_T = 0.59$, our guess for r_C is about 0.49. This yields $\tilde{k} = 0.59/0.49 = 1.20$, which is less than $k^* = 1.25$ and is therefore consistent with a true negative small effect ($\tilde{\Delta} = 0.59 \cdot (0.50 - 0.625/1.20) = -0.012$).

Thus, in addition to allow researchers to explore how severe the pattern of differential registration between the treated and control groups would have to be to reverse the sign of the estimated difference in turnout-to-registration rates, our simple sensitivity analysis also allows researchers to

explore whether the estimated effect is robust to a particular differential rate of registration based on prior knowledge.

Application: The effects of election eligibility on subsequent turnout

We now illustrate the problem of differential registration bias and our approach with two empirical studies of political socialization that focus on the relationship between voting eligibility and subsequent voter turnout. In Study 1, we present an original analysis of voter file data from 42 states. In Study 2, we replicate results in a recent article by Holbein and Hillygus (2015). In both cases, the empirical strategy is based on a comparison of voters who were born close to an election eligibility cutoff, which we describe in more detail below.

Research in political socialization has found long-lasting effects of early experiences and events like parent socialization (Niemi and Jennings 1991; Jennings, Stoker and Bowers 2009) and draft status during the Vietnam War (Erikson and Stoker 2011). The most common and important socializing events for many people as they approach or enter adulthood are elections—the time when politics is most salient in national life. Sears and Valentino (1997), for instance, find that presidential elections appear to be especially potent in forming the political views of adolescents.

These questions have been re-engaged by an emerging literature that studies the effects of initial election eligibility on voter turnout and other political behaviors using a quasi-experimental approach (e.g., Coppock and Green 2015; Dinas 2012, 2014; Holbein and Hillygus 2015; Meredith 2009; Mullainathan and Washington 2009). By comparing later turnout and political attitudes among a group of voters whose 18th birthday fell very close to a general election, these studies seek to leverage as-if random variation in birth timing to compare individuals who had the opportunity to take part in an election and those who did not but are assumed to be otherwise identical. This research strategy, which is based on the discontinuous change in voting eligibility that occurs at election day, is an application of a regression discontinuity (RD) design in which a treatment is given on the basis of whether a score exceeds a known cutoff.

These RD-style designs are a promising approach to estimating election socialization effects because they can be implemented at a relatively low cost using voter files. Of course, the validity of

the results depends on the assumption that the eligibility treatment was randomly or as-if randomly assigned.⁴ While the plausibility of this assumption will of course vary depending on context, our focus in this article is not on RD design assumptions. We use RD-based applications simply as illustrations; our approach is general and applies to all non-experimental turnout studies, not just RD designs.

Studying eligibility effects with a regression discontinuity design

The two examples we analyze below are based on an RD design of the type described above. We therefore discuss this design in more detail before turning to the presentation of each study. The defining feature of an RD design is that subjects are assigned a score and receive treatment if their score exceeds a known cutoff. In the U.S., a discontinuity in voting eligibility occurs when citizens turn eighteen years of age. As a result of the 26th Amendment to the U.S. Constitution (adopted in 1971), people who turn eighteen on or before election day can cast a vote but those who will turn eighteen after election day are ineligible to vote. Thus, date of birth exactly determines voting eligibility and an RD design can be used to study the effects of eligibility on turnout.

An important feature distinguishing RD designs based on date of birth from most uses of RD is that the score that determines treatment, birthdate, is a discrete variable, which invalidates most identification and estimation results in the RD literature. To address this issue, we adopt the framework in Cattaneo, Frandsen and Titiunik (2015), which analyzes the RD design as a local randomized experiment in a fixed window around the cutoff and does not require a continuous running variable.⁵ In our context, this randomization-based RD approach entails assuming that voting eligibility is as-if randomly assigned for people with birthdays near election day. Since the number of observations in our applications is large, we do not use the randomization inference methods discussed in Cattaneo, Frandsen and Titiunik (2015). All our inferences are based in large-sample approximations; we report confidence intervals for the difference-in-means between just-eligible and just-ineligibles based on Wald tests—we use t-tests for our turnout-to-registration analysis and

⁴Several of these studies test a habit-formation hypothesis by assuming election eligibility affects subsequent voting only through its effect on current voting (Meredith 2009; Dinas 2012; Coppock and Green 2015). We do not consider the validity of this assumption and instead focus on the direct effect of eligibility on subsequent voting.

⁵See Cattaneo, Titiunik and Vazquez-Bare (2015) for a comparison of this randomization-based RD approach to the more standard continuity-based approach.

proportions tests (Newcombe 1998) when we consider turnout as proportion of births.

In order to adopt this local experiment framework, it is crucial to focus on individuals who are born close in time: although factors such as education and civic engagement may differ between 17-year-olds and 19-year-olds, we do not generally expect major differences between someone who is exactly 18 on election day and someone born just a few days later. Thus, in Study 1, we focus our analysis on individuals who turn eighteen within eight days of election day and assume that eligibility can be considered as-if randomly assigned between those individuals born on election day or the three days earlier (the treatment group) and those born one to four days later (the control group). In Study 2, we use a wider window around election day to ensure comparability with the approach used in Holbein and Hillygus (2015).

Both studies estimate the effects of voting eligibility on subsequent turnout using voter file data, which means they are vulnerable to differential registration bias. For example, if just-eligibles are more likely to be registered than just-ineligibles due to the longer period in which they could participate in the political process or be mobilized by campaigns, we could observe a spurious differential in turnout-to-registration rates even if there is no difference in the underlying true turnout rates between the two populations.

Study 1: Voter eligibility effects in Catalist data

Our first study is an original application that investigates the effects of voting eligibility on subsequent turnout with an RD design based only on the closest observations to the election day cutoff. Specifically, we examine three cohorts who were narrowly (in)eligible to vote in the 2004, 2006, and 2008 elections, considering only those registrants born within just four days of the election eligibility cutoff—a far narrower window than previous studies, which have used windows measured in months (Meredith 2009; Dinas 2012, 2014; Holbein and Hillygus 2015) or years (Mullainathan and Washington 2009; Coppock and Green 2015).

In addition, our sample is comprehensive. Our data are drawn from voter files in 42 U.S. states and the District of Columbia and include eligibility variation and turnout data from several national elections. Our data source is voter registration files that were collected, cleaned, and supplemented

by the private company Catalist.⁶ We collected a random sample of voters in the Catalist file born in the eight days around the cutoff date for being eligible to vote (i.e., for being 18 years old on or before election day) in the 2004, 2006, and 2008 elections.⁷

The three cohorts of individuals in our data were born in 1986, 1988, and 1990, respectively. For example, the 1990 cohort treatment group was born from November 1–4 and were thus eighteen years old and eligible to vote on November 4, 2008, while the control group was born from November 5–8, 1990. Unfortunately, Catalist’s data on unregistered voters are sparse and unreliable, which forces us to focus—like other analysts—on the universe of registrants and thereby introduces the possibility of differential registration bias. The final dataset includes a total of 49,271 observations in our target windows among the three birth cohorts.⁸

Effects of eligibility on turnout-to-registration rates

Before turning to our results, Table 3 explains how we present the findings from our analysis. We compare the behavior of the treatment group of just-eligibles—those who were born just before the election eligibility cutoff—with the control group of just-ineligible voters born just after the cutoff in later elections. The election in the year the cohort turned 18 is denoted E1 and subsequent elections are denoted E2, E3, and E4. For instance, E1 for the 1986 cohort is the 2004 election and the 2006, 2008, and 2010 elections are E2, E3, and E4, respectively, for that cohort.

We analyze the effect of election eligibility in the year voters turned 18 on subsequent turnout-to-registration rates in Table 4, which compares just-eligible and just-ineligible voters who were

⁶Colorado, Massachusetts, New Jersey, Oklahoma, South Carolina, Vermont, and Washington were excluded due to school entry cutoff dates that overlapped with the election eligibility window, creating potential discontinuous differences in education levels. Illinois was excluded due to legal restrictions on state voter file use.

⁷Ansolahehere and Hersh (2010) use Catalist data to analyze the quality of state voter files and find that “Identifying information such as birthdates are generally well collected.” They do identify some problems with missing birth dates and unusual concentrations of voters with particular birth dates but these should not affect the validity of our design. (The only unusual date within our window is November 11, which they find to be unusually prevalent in Texas, but we observe no evidence of a problem in our data [results available upon request].) Catalist not only cleans and processes data from state voter files, which includes tracking individuals who move between states and/or are purged from voter files, but fills in exact birth dates from commercial sources when possible for states that only release month of birth, allowing us to use exact birthdates even in states that do not release them.

⁸We drop all observations missing exact birthdates, those with birthdates outside the target range, and those recorded as voting in elections for which they should have been ineligible given their reported birthdate. See the Supporting Information for details on the number of excluded observations.

Table 3: Birth years and election years in 2011 Catalist data

Year	E1	E2	E3	E4
1986	2004	2006	2008	2010
1988	2006	2008	2010	-
1990	2008	2010	-	-

Table 4: Turnout-to-registration rates by voting eligibility

A. 1986 cohort (first election for just-eligibles: 2004 presidential)			
	E2 (2006 midterm)	E3 (2008 presidential)	E4 (2010 midterm)
Eligibility effect	2.12 [1.07, 3.17]	-2.16 [-3.61, -0.72]	-2.56 [-3.70, -1.43]
Control group	14.59	52.18	20.37
B. 1988 cohort (first election for just-eligibles: 2006 midterm)			
	E2 (2008 presidential)	E3 (2010 midterm)	E4 (2012 presidential)
Eligibility effect	-1.50 [-2.99, -0.01]	-1.91 [-3.08, -0.73]	- -
Control group	55.51	20.11	
C. 1990 cohort (first election for just-eligibles: 2008 presidential)			
	E2 (2010 midterm)	E3 (2012 presidential)	E4 (2014 midterm)
Eligibility effect	-3.06 [-4.45, -1.66]	- -	- -
Control group	23.10	-	-

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792). Brackets show 95% confidence intervals based on a differences-in-means Wald test.

born in the week surrounding the eligibility cutoff.⁹ These findings initially seem to contradict findings that eligibility increases subsequent turnout (e.g., Meredith 2009; Dinas 2012; Coppock and Green 2015). While we find a significant positive effect of eligibility on turnout-to-registration rates for the 1986 cohort in the 2006 election, the estimated effect is *negative* and significant for the 1986 cohort in the 2008 and 2010 elections, the 1988 cohort in the 2008 and 2010 elections, and the 1990 cohort in the 2010 election.

Specifically, registered voters who were born in 1986 and were just eligible to vote in 2004 were significantly more likely to turn out in 2006 than those who were just ineligible. The estimated treatment effect is approximately two percentage points (2.12 percentage points, 95% CI: 1.07,

⁹Balance tests in the Supporting Information show little imbalance on observables among registrants in our data.

Table 5: Sensitivity analysis

A. 1986 cohort (first election for just-eligibles: 2004 presidential)					
E2 (2006 midterm)		E3 (2008 presidential)		E4 (2010 midterm)	
$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*
2.12	0.87	-2.16	1.04	-2.56	1.14
B. 1988 cohort (first election for just-eligibles: 2006 midterm)					
E2 (2008 presidential)		E3 (2010 midterm)		E4 (2012 presidential)	
$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*
-1.50	1.03	-1.91	1.05	-	-
C. 1990 cohort (first election for just-eligibles: 2008 presidential)					
E2 (2010 midterm)		E3 (2012 presidential)		E4 (2014 midterm)	
$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*
-3.06	1.15	-	-	-	-

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792).

3.17), which is a substantial increase relative to the low baseline turnout rate for young voters in a midterm election but relatively modest in absolute terms. However, this apparent effect reverses by the second and third subsequent elections—just-eligible voters born in 1986 were significantly *less* likely to vote in 2008 and 2010 than their just-ineligible counterparts among the registered voters in our data. We find a similar *negative* relationship between eligibility and subsequent turnout-to-registration rates for just-eligible registered voters born in 1988 cohort in 2008 and 2010 and the 1990 cohort in 2010. (A full set of RD plots illustrating these estimates is included in the Supporting Information.)

Sensitivity analysis: Assessing differential registration scenarios

We now conduct a sensitivity analysis to diagnose how robust the results above are to differential registration. Table 5 presents the sensitivity tests for our data. Again, the key term is k^* , which is the ratio of registration between the treatment and control groups that would produce the observed difference in turnout-to-registration rates if the true difference in the turnout-to-population rates were zero. When k^* is close to 1, the result is highly sensitive to a small difference in registration rates. By contrast, a k^* that is farther from 1 indicates that the result is more robust.

These results indicate that the positive eligibility effect we observe for turnout-to-registration rates in the 2006 election among the 1986 cohort appears to be robust. The estimated value of k^* is 0.87, which means that just-eligibles would have to register at a *lower* rate than just-ineligibles to explain the result if the true effect on turnout-to-population rates was zero. In the absence of pre-registration laws, it is plausible to assume that just-eligible voters are more likely to register than just-ineligible voters so we can therefore place some confidence in this result.

By contrast, the other estimated values of k^* suggest that the negative estimates of the effect of eligibility on subsequent turnout-to-registration rates in Table 5 are highly sensitive to differential registration. These estimates have corresponding k^* values of 1.03–1.15, which means that only slight registration differentials in the expected direction (i.e., $r_T > r_C$) could produce the observed negative turnout-to-registration effects. Correspondingly, if the registration differentials were larger than k^* , the true effect on the turnout-to-population rate would be positive.

Assessing k^* using external data

The values of k^* reported above indicate that relatively small differences in registration rates between the treatment and control groups could explain the observed negative results. We now use birth totals as a proxy for the VEP to briefly explore whether differences of these magnitudes are plausible in this application.

We calculate daily birth totals within the eight-day window around election day in the 1986 and 1988 cohorts for our sample of 42 states and the District of Columbia using data from *Vital Statistics of the United States*. Exact birth dates were redacted from these data starting in 1989, preventing us from constructing similar estimates for the 1990 cohort. We therefore instead estimate daily birth totals for our sample states by scaling total U.S. births for each birthdate in our window from the 1990 edition of *Vital Statistics* by the proportion of the population living in those states at the time.¹⁰

Using these data, we divide the total number of registrants in the treatment and control groups by birth totals from public data sources, which produces estimates of r_T and r_C . These figures

¹⁰The proportion of the U.S. population living in the states in our sample was stable during this period so we did not further adjust these estimates to account for interstate migration.

Table 6: Registration rates as of 2011 by voting eligibility as proportion of births

Year	Treated	Control	\hat{k}	Sensitivity (k^*)		
				E2	E3	E4
1986	28.42	25.99	1.09	0.87	1.04	1.14
1988	24.47	23.41	1.04	1.03	1.05	-
1990	22.59	14.48	1.56	1.15	-	-

are *not* valid estimates of registration rates because our data are a random sample from Catalist’s voter file and do not include every voter registered on the dates in question in our sample states. However, the *difference* between these estimated registration rates is a valid estimate of differential registration bias in our window around election day due to the use of random sampling.

We report the estimated registration rates for the treatment and control groups in Table 6 as well as estimates of the differential registration factor \hat{k} . The registration rate as a proportion of births is much higher in the treatment than in the control group in each row (all $p < .01$). These differences are greatest for the 1990 cohort because just-ineligibles has had less time to “catch up” to just-eligibles, but persist even among the 1986 cohort seven years after turning 18.

Most notably, our estimates of the differential registration factor \hat{k} are well within the range that the sensitivity analysis in Table 5 suggests could explain our negative turnout-to-registration results. For the 1986 cohort, \hat{k} is 1.09 and the values of k^* that could explain the negative turnout-to-registration estimates in E3 and E4 are, respectively, 1.04 and 1.14. Likewise, \hat{k} is 1.04 for the 1988 cohort and the E2 and E3 values of k^* are, respectively, 1.03 and 1.05. Finally, the \hat{k} value of 1.56 for the 1990 cohort greatly exceeds the 1.15 estimate of k^* for E2.¹¹

Another way to look at these findings is to perform a second RD analysis comparing turnout rates between just-eligible and just-ineligible voters using birth totals rather than registrants as the denominator (corresponding RD plots are provided in the Supporting Information). Consistent with the results in Table 6, Table 7 shows that when we use birth totals in the denominator, the results are largely the opposite of what we found when we conditioned on registration. In some elections we consider, being eligible to vote at age 18 significantly increases downstream turnout-to-population rates (E2 and E3 for the 1986 cohort and E2 for the 1990 cohort), which is consistent

¹¹Because we observe registration only in 2011, our estimate \hat{k} is constant within each birth cohort.

Table 7: Turnout rates by voting eligibility as a proportion of births

A. 1986 cohort (first election for just-eligibles: 2004 presidential)			
	E2 (2006 midterm)	E3 (2008 presidential)	E4 (2010 midterm)
Eligibility effect	0.96	0.65	-0.23
	[0.65, 1.27]	[0.13, 1.18]	[-0.57, 0.10]
Control group	3.79	13.56	5.29
B. 1988 cohort (first election for just-eligibles: 2006 midterm)			
	E2 (2008 presidential)	E3 (2010 midterm)	E4 (2012 presidential)
Eligibility effect	0.22	-0.25	-
	[-0.28, 0.72]	[-0.56, 0.06]	-
Control group	12.99	4.70	
C. 1990 cohort (first election for just-eligibles: 2008 presidential)			
	E2 (2010 midterm)	E3 (2012 presidential)	E4 (2014 midterm)
Eligibility effect	1.18	-	-
	[0.90, 1.47]	-	-
Control group	3.34	-	-

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792). Brackets show 95% confidence intervals based on a differences-in-proportions Wald test.


with prior findings (Meredith 2009; Dinas 2012; Coppock and Green 2015). In the other elections, the estimated effect of eligibility is null.¹²

The reversal of the negative effects on turnout-to-registration rates in Table 7 is the result by differences in birth counts between groups. Figure 3 illustrates the phenomenon using data for the 1986 cohort. Even though the total registration and vote counts are similar between groups, birth counts are higher in the control group, considerably reducing the turnout-to-population rates relative to the treatment group. The higher number of births in the control group is the result of the well-known pattern of day-level variation in birth rates. As we show in the Supporting Information, the treatment windows of four days in the 1986, 1988, and 1990 cohorts all include two weekend days, when birth rates are typically lower in the U.S., while the control windows include only weekdays. The control group will therefore be larger than the treatment group and as a result turnout rates calculated as a proportion of births will fall disproportionately more for the control group.

This difference in the density of observations around the cutoff does not necessarily invalidate

¹²As we show in the Supporting Information, our results are unchanged when we exclude states with preregistration.

Figure 3: Total population, registration, and voters for 1986 cohort



Note: Our data are a random sample from Catalist’s voter file and therefore underestimate the turnout-to-population rates for both the treatment and control groups. However, because the data were drawn randomly, we can accurately estimate the difference in turnout-to-population rates between groups. Voting is measured in the 2008 election, and registration is measured in 2011.

the RD design, which would only occur if the likelihood of having a weekend birth is correlated with potential voting behavior. However, our analysis shows that small seasonal variations in births can reverse observed turnout-to-registration effects.

Study 2: Preregistration effects in Florida

Our second study is based on the recent work by Holbein and Hillygus (2015), who investigate the effects of preregistration on future turnout among young people. Preregistration laws typically allow voting-ineligible 16-year-old or 17-year-old citizens to complete a registration application so that they are automatically added to the registration rolls once they turn eighteen and become eligible to vote. The authors present analyses of both cross-state data from the Current Population Survey and the Florida voter file. In each case, they find evidence that the availability of preregistration has a positive effect on young people’s subsequent turnout, *increasing* the probability that people who are narrowly ineligible will vote in future elections.

We choose to reanalyze the Holbein and Hillygus (2015) study for two reasons. First, the results they present contrast strongly with Coppock and Green (2015), who show, as we do in Study 1, that just-eligible voters are more likely to vote in subsequent elections than just-ineligibles. We believe our sensitivity approach can contribute to reconciling the conflicting findings in these papers. Second, this case provides another opportunity to demonstrate the importance of sensitivity analysis in evaluating claims about turnout made using voter file data when external data sources face severe limitations—possibly more severe than in Study 1. In particular, we are only able to construct an approximate measure of the total population born on the dates being studied who live in Florida due to limitations on data availability as well as patterns of migration into Florida.

We focus exclusively on Holbein and Hillygus’s (2015) second analysis, which compares voter turnout among narrowly eligible and narrowly ineligible Florida voters who were born in 1990 close to the voting-eligibility cutoff for the 2008 presidential election. Holbein and Hillygus (2015) use this design to estimate the effects of preregistration. In Florida, where preregistration is allowed, narrowly ineligible voters are exposed to the opportunity to preregister, while most of those who are narrowly eligible to vote register “regularly” (i.e., when they are already eighteen). They conceptualize narrowly ineligible voters as the treatment group and narrowly eligible voters as the control group; ineligibility is an instrument for preregistration, which is the treatment of interest. Their analysis is based on a fuzzy RD design where ineligibility induces preregistration.

Our re-analysis of Holbein and Hillygus’ Florida results, which uses the comprehensive replication materials they generously provided, differs from their original study in important ways. We are primarily interested in illustrating how differential registration patterns between treatment and control groups can affect turnout studies that calculate turnout rates as a proportion of registration. For this reason, we re-analyze Holbein and Hillygus (2015) as a sharp regression discontinuity design where, as in our Study 1, the treatment of interest is voting eligibility (as opposed to preregistration), narrowly eligible voters are the treatment group, and narrowly ineligible voters are the control group. Our design is thus analogous to the intent-to-treat (ITT) analysis that they report in the article except that the treatment and control group labels are inverted.

Table 8: Eligibility Effects on Turnout-to-Registration—1990 Florida data

Window	T_T^{Reg} (turnout-to-reg eligibles)	T_C^{Reg} (turnout-to-reg ineligibles)	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	k^*
± One month	49.08	51.88	-2.80 [-3.92, -1.68]	1.06
± Two months	49.28	51.83	-2.55 [-3.33, -1.77]	1.05

Note: Source is Florida voter file data for 1990 births one month before or after November 4th; sample size = 30,979. Brackets show 95% confidence intervals based on a differences-in-means Wald test.

Effects of eligibility on turnout-to-registration rates

We first estimate the effect of voting eligibility on future turnout-to-registration rates and then conduct a sensitivity analysis to determine if the results could be driven by differential registration bias. For our main analysis, we subset the data to people born from October 4–December 4, 1990 to match the Holbein and Hillygus (2015) window of approximately one month on either side of election day. Within this window, we treat the assignment of voting eligibility in 2008 as locally random and compare the turnout-to-registration rate in 2012 between just-eligibles and just-ineligibles. We also consider a larger window of two months on either side of the cutoff.

Table 8 reports the results for both windows. In the one-month window, our estimated treatment effect on turnout-to-registration rates is -2.80 percentage points, meaning that just-ineligible registrants who were exposed to the option to preregister in 2008 voted at a higher rate in 2012 (51.88%) than registrants who in 2008 were just-eligible (49.08%). This estimate is very close to the three percentage-point effect that Holbein and Hillygus (2015) report for their ITT estimate.¹³ In the larger window of four months surrounding the cutoff, we find a similar pattern, with just-ineligibles again being slightly more likely to turn out than just-eligibles as a proportion of registrants. Both negative effects are significantly different from zero at 5% level.

¹³The difference is likely due to the fact that, unlike Holbein and Hillygus (2015), our analysis reports a simple difference in means and does not include controls.

Sensitivity analysis: Assessing differential registration scenarios

Our goal is to establish the robustness of the above finding to potential differential registration. As shown in the last column of Table 8, our sensitivity analysis gives a k^* value of 1.06 for the one-month window used in Holbein and Hillygus (2015), which means that if the rate of registration were six percent higher in the treatment than in the control group, the negative turnout-to-registration effect we observe would reflect a null turnout-to-population effect (and any difference greater than six percent would switch the effect from negative to positive). In the two-month window, the sensitivity of the results is even larger, with a k^* equal to 1.05.

In 2008, the nationwide percentage of 18-year-olds who reported being registered to vote was approximately 49% (Herman and Forbes 2010). Assuming that the registration rate among just-eligible (treated) Florida voters within the window we consider was also 49%, a k^* value of 1.06 implies that the rate of registration among just-ineligible (control) voters would only have to be 46.23% to change the sign of the point estimate—a difference of just 2.77 percentage points. We therefore conclude that, as in Study 1, the negative eligibility effect found in Florida is not robust to *positive* differential registration—i.e., to a situation where the registration rate is higher in the treatment than in the control group.¹⁴

Assessing k^* empirically using external data

As in Study 1, we now use births in 1990 to provide approximate estimates of the rate of differential registration. Since the CDC's National Vital Statistics do not contain birth rates disaggregated by state and exact date in 1990 or later years, we cannot use an exact one-month window around the cutoff date of November 4, 1990. Instead, we approximate the correct birth totals using the October totals for the treatment group and the November totals for the control group. Again, this approximate analysis is intended to provide the best available evidence about plausible values that k can take in this application.

In Table 9, we compare the results using both registered voters and births as the denominator

¹⁴The main focus of Holbein and Hillygus (2015) is not the ITT, but rather the TOT effect. This quantity is the ITT effect divided by the rate of preregistration in the just-ineligible group (preregistration is zero by construction among eligibles). A change in the sign of the ITT effect would thus result in a change in sign in the TOT effect as well.

Table 9: RD estimates of eligibility effect in 1990 Florida cohort

Denominator	T_T^{Reg} (turnout eligibles)	T_C^{Reg} (turnout ineligible)	$T_T^{\text{Reg}} - T_C^{\text{Reg}}$	\hat{k}
Two-month window				
Registration	49.08	51.88	-2.80 [-3.92, -1.68]	-
Births	44.81	47.48	-2.67 [-3.74, -1.60]	0.94
Four-month window				
Registration	49.29	51.84	-2.55 [-3.33, -1.77]	-
Births	48.76	45.50	3.25 [2.61, 4.10]	1.13

Note: Florida voter file data for 1990 births within one/two month(s) of November 4th. One-month window sample size = 30,979. Two-month window sample size = 64,286. Estimated births calculated from Florida vital records for October and November 1990. Brackets show 95% confidence intervals based on a differences-in-proportions Wald test.

in the two windows, reproducing the effects reported in Table 8 for easy comparability. In the one-month window, the effect remains negative when we calculate turnout based on approximate births in the window. Indeed, the point estimate is very similar to the effect based on registration: we estimate that 2012 turnout among Floridians who were narrowly eligible in 2008 was -2.67 percentage points lower than their just-ineligible counterparts.

However, this finding changes when we consider a larger window of two months on either side of the cutoff. Table 9 shows that the turnout-to-registration effect estimate in this window continues to resemble the one found by Holbein and Hillygus (2015). However, when we calculate the estimated turnout-to-births rate in this window, the direction of the effect flips and the 2012 turnout rate among just-eligibles is found to be *higher* than among just-ineligibles.

This result can be explained by the fact that the number of births in Florida during November 1990 was significantly lower than in October. Except for November, total monthly births in Florida between July and December 1990 were above 17,000. The total number of births in November 1990, however, was only 16,289. Thus, although the number of people who voted in 2012 was very similar in the treatment and the control groups (7,888 and 7,734, respectively), the control

group's approximate size was apparently much smaller. We show in the Supporting Information that there is significant seasonal variation in monthly births in the U.S., which again highlights the difficulty of accurately estimating the relevant VEP and the necessity of assessing the sensitivity of a study's conclusions to departures from accurate estimation of the VEP.¹⁵

Despite those caveats, our re-analysis suggests that the Holbein and Hillygus (2015) findings are not robust. Our estimated value of k^* indicates that the Holbein and Hillygus (2015) finding would become null in the presence of small amounts of differential registration bias. Using birth counts as a tentative approximation of the VEP, we also find conflicting results when we change the window around Election Day. By contrast, both Study 1 and Coppock and Green (2015) provide strong evidence that just-eligibles turn out more, not less.

Conclusion

Research using voter file data is more common than ever in American politics. Frequently, scholars use voter files to investigate the effect of a non-experimental treatment on turnout (or partisan registration). However, the voting-eligible population (VEP), which is necessary to calculate the turnout effects of the treatment of interest, is unavailable. Some researchers simply use the population of registrants as the universe of analysis, an approach that implicitly conditions on voter registration and risks what we call differential registration bias, potentially distorting estimates of how a treatment affects turnout rates. Others choose to study only the subset of the population that was registered before the intervention of interest occurred or instead approximate the VEP from secondary sources.

We study formally the problem of differential registration and develop a new approach to sensitivity analysis that allows scholars to assess how robust their estimates are to potential differences in registration rates between treatment and control groups. Our approach is most helpful when only a cross-section of registered citizens is available, but is also useful as a complementary analysis when only a subset of the registered population is studied before and after the treatment or when a secondary source is used as denominator.

¹⁵Another concern is that birth data do not account for migration to and from Florida during the 1990–2008 period. These population shifts mean that birth counts are an imprecise measure of VEP.

We illustrate the use of these methods with two studies of turnout based on voter file data—an original analysis of voter file data from 42 U.S. states and a reanalysis of Holbein and Hillygus (2015) using data from Florida voter files. In both cases, we show that comparisons of turnout-to-registration rates are highly sensitive to differential registration bias and would reverse if the registration rate were slightly higher in the treatment group than the control group. Moreover, when we use birth counts to approximate the total population in each group, we find that the estimated differential registration factor is very close to the value that would reverse the sign of the estimated turnout-to-registration effects.

These findings illustrate the complex tradeoffs associated with changes in the data used in studies of voter turnout and political behavior more generally. Until recently, turnout studies often relied on survey data, which suffers from two primary limitations: measurement error in self-reports, especially overreporting (e.g., Ansolabehere and Hersh 2012; Fraga 2015); and missing data due to nonresponse bias, which overrepresents individuals who are more knowledgeable and interested in politics (e.g. Couper 1997). As we show, studies that rely on voter files eliminate the overreporting problem but not the inferential challenge posed by missing data, which now consists instead of those citizens who are not registered rather than those who do not take part in surveys.

Thankfully, though, we show that the missing data problem is less severe for voter file data than in surveys because, by construction, all eligible voters who are not registered did not cast a vote. (In contrast, survey nonrespondents may or may not have voted.) Our sensitivity analysis directly incorporates this additional information about nonregistrants, leaving only one unknown variable (the differential registration factor k) to be varied. By contrast, though a similar sensitivity analysis could be used to assess nonresponse bias in surveys, the fact that the outcome of nonrespondents is completely unknown implies that such an analysis would have to be based on standard partial identification results for average treatment effects (e.g. Manski 2003), leading to a much wider range of possible effects and decreasing its appeal for practitioners.

Our study also offers important methodological lessons. The crucial property on which our approach is based—the size of the treatment and control populations are not known but the values of the missing outcomes are known with certainty—is most common in non-experimental studies of voter turnout, but may also occur in other literatures such as studies of how often the Supreme

Court reverses lower court decisions¹⁶ or whether militarized disputes between countries escalate into war.¹⁷

In addition, our study illustrates that even rigorous designs can be vulnerable to post-treatment bias despite other assumptions being met. If we condition on variables affected by the treatment of interest in any way, our estimates of causal effects may be severely misleading. This possibility is especially insidious in cases where the treatment affects the composition of the sample studied. Differential registration bias is only one way in which failing to take this threat into account can lead us to mistaken inferences about how politics works.

References

- Ansolabehere, Stephen and Eitan Hersh. 2010. “The Quality of Voter Registration Records: A State-by-State Analysis.” Report, Harvard University. Downloaded August 22, 2015 from http://www.eitanhersh.com/uploads/7/9/7/5/7975685/reg_quality_report_8-5-10.pdf.
- Ansolabehere, Stephen and Eitan Hersh. 2012. “Validation: What big data reveal about survey misreporting and the real electorate.” *Political Analysis* 20(4):437–459.
- Barber, Michael and Kosuke Imai. 2014. “Estimating Neighborhood Effects on Turnout from Geocoded Voter Registration Records.” *Working Paper, Princeton University* .
- Barreto, Matt. A. 2007. “Sí Se Puede! Latino Candidates and the Mobilization of Latino Voters.” *American Political Science Review* 101(3):425–441.
- Barreto, Matt A, Gary M Segura and Nathan D Woods. 2004. “The Mobilizing Effect of Majority-Minority Districts on Latino Turnout.” *American Political Science Review* 98(01):65–75.
- Cattaneo, Matias D., Brigham Frandsen and Rocío Titiunik. 2015. “Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate.” *Journal of Causal Inference* 3(1):1–24.
- Cattaneo, Matias D., Rocio Titiunik and Gonzalo Vazquez-Bare. 2015. “Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality.” working paper, University of Michigan .

¹⁶The cases heard by the Supreme Court are of course a selected sample drawn not just from the appeals it receives but the set of *potential* cases that could be filed. If a treatment influences the likelihood of a case being filed, scholars could observe a seemingly paradoxical result of the sort identified by McGuire, Vanberg, Smith and Caldeira (2009).

¹⁷As Reed (2000) notes, many analysts have analyzed the likelihood of escalation among observed disputes, neglecting how predictors of interest may also influence the likelihood of dispute initiation among the unknown set of *potential* disputes that could be initiated.

- Cepaluni, Gabriel and F. Daniel Hidalgo. 2016. "Compulsory Voting Can Increase Political Inequality: Evidence from Brazil." *Political Analysis, forthcoming* .
- Citrin, Jack, Donald P Green and Morris Levy. 2014. "The effects of voter ID notification on voter turnout: Results from a large-scale field experiment." *Election Law Journal* 13(2):228–242.
- Coppock, Alexander and Donald Green. 2015. "Is Voting Habit Forming? New Evidence from Experiments and Regression Discontinuities." *American Journal of Political Science* .
- Couper, Mick P. 1997. "Survey Introductions and Data Quality." *Public Opinion Quarterly* pp. 317–338.
- Dinas, Elias. 2012. "The formation of voting habits." *Journal of Elections, Public Opinion & Parties* 22(4):431–456.
- Dinas, Elias. 2014. "Does choice bring loyalty? Electoral participation and the development of party identification." *American Journal of Political Science* 58(2):449–465.
- Elwert, Felix and Christopher Winship. 2014. "Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable." *Annual Review of Sociology* 40(31–53).
- Enos, Ryan D. 2016. "What the Demolition of Public Housing Teaches Us about the Impact of Racial Threat on Political Behavior." *American Journal of Political Science, forthcoming* .
URL: <http://dx.doi.org/10.1111/ajps.12156>
- Erikson, Robert S. and Laura Stoker. 2011. "Caught in the draft: The effects of Vietnam draft lottery status on political attitudes." *American Political Science Review* 105(2):221–237.
- Fowler, James H., Laura A. Baker and Christopher T. Dawes. 2008. "Genetic variation in political participation." *American Political Science Review* 102(2):233–248.
- Fraga, Bernard L. 2015. "Candidates or Districts? Reevaluating the Role of Race in Voter Turnout." *American Journal of Political Science, forthcoming* .
URL: <http://dx.doi.org/10.1111/ajps.12172>
- Fraga, Bernard L. 2016. "Redistricting and the Causal Impact of Race on Voter Turnout." *The Journal of Politics* 78(1):19–34.
- Fraga, Bernard L. and Julie Lee Merseth. 2016. "Examining the Causal Impact of the Voting Rights Act Language Minority Provisions." *Journal of Race, Ethnicity, and Politics, forthcoming* .
- Gerber, Alan S., Donald P. Green and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment." *American Political Science Review* 102(1):33–48.
- Herman, Jody and Lauren Forbes. 2010. "Engaging America's Youth through High School Voter Registration Programs." Project Vote. Downloaded August 22, 2015 from http://funderscommittee.org/files/high_school_voter_registration_programs.pdf.

- Hersh, Eitan D. 2013. "Long-term Effect of September 11 on the Political Behavior of Victims' Families and Neighbors." *Proceedings of the National Academy of Sciences* 110(52):20959–20963.
- Hersh, Eitan D. 2015. *Hacking the Electorate: How Campaigns Perceive Voters*. Cambridge University Press.
- Hersh, Eitan D. and Clayton Nall. 2016. "The Primacy of Race in the Geography of Income-Based Voting: New Evidence from Public Voting Records." *American Journal of Political Science*, forthcoming .
URL: <http://dx.doi.org/10.1111/ajps.12179>
- Holbein, John B. and D. Sunshine Hillygus. 2015. "Making Young Voters: The Impact of Preregistration on Youth Turnout." *American Journal of Political Science* Early View.
- Jennings, M. Kent, Laura Stoker and Jake Bowers. 2009. "Politics across generations: Family transmission reexamined." *Journal of Politics* 71(3):782–799.
- Manski, Charles. 2003. *Partial Identification of Probability Distributions*. New York: Springer-Verlag.
- McDonald, Michael P. and Samuel L. Popkin. 2001. "The Myth of the Vanishing Voter." *American Political Science Review* 95(4):963–974.
- McGuire, Kevin T., George Vanberg, Charles E. Smith and Gregory A. Caldeira. 2009. "Measuring policy content on the U.S. Supreme Court." *Journal of Politics* 71(4):1305–1321.
- Meredith, Marc. 2009. "Persistence in political participation." *Quarterly Journal of Political Science* 4(3):187–209.
- Mullainathan, Sendhil and Ebonya Washington. 2009. "Sticking with Your Vote: Cognitive Dissonance and Political Attitudes." *American Economic Journal: Applied Economics* 1(1):86–111.
- Newcombe, Robert G. 1998. "Interval estimation for the difference between independent proportions: comparison of eleven methods." *Statistics in medicine* 17(8):873–890.
- Niemi, Richard G. and M. Kent Jennings. 1991. "Issues and inheritance in the formation of party identification." *American Journal of Political Science* 35(4):970–988.
- Reed, William. 2000. "A unified statistical model of conflict onset and escalation." *American Journal of Political Science* 44(1):84–93.
- Rosenbaum, Paul R. 1984. "The Consequences of Adjustment for a Concomitant Variable That Has Been Affected by the Treatment." *Journal of the Royal Statistical Society, Series A (General)* 147(5):656–666.
- Sears, David O. and Nicholas A. Valentino. 1997. "Politics matters: Political events as catalysts for preadult socialization." *American Political Science Review* 91(1):45–65.

Supporting Information for “Differential Registration Bias in Voter File Data: A Sensitivity Analysis Approach”

Brendan Nyhan* Christopher Skovron† Rocío Titiunik‡

January 12, 2016

Contents

S1 Excluded Catalist data in Study 1	2
S2 Balance checks for Study 1	3
S3 RD plots for Study 1	5
S3.1 Turnout-to-registration rates	5
S3.2 Turnout-to-births rates	7
S4 Results for Study 1 in states without preregistration	9
S5 Total births by day and month	10

*Assistant Professor, Department of Government, Dartmouth College, nyhan@dartmouth.edu

†Ph.D. Candidate, Department of Political Science, University of Michigan, cskovron@umich.edu.

‡Assistant Professor, Department of Political Science, University of Michigan, titiunik@umich.edu.

S1 Excluded Catalist data in Study 1

As mentioned in the paper, we excluded some individuals who were in the raw data provided to us by Catalist. Table S1 describes the observations excluded by category.

Table S1: Observations excluded from analysis in study 1

Description	Observations
Original dataset	57,031
After dropping missing exact birthdates	54,332
After dropping outside birth targets	51,705
After dropping those recorded as voting when they should have been ineligible	51,472
After dropping those with no registration year listed	49,271

S2 Balance checks for Study 1

To demonstrate that the as-if randomized assumption of our RD design is plausible among the registrants in our data for Study 1, we compare the demographic characteristics of just-eligible and just-ineligible voters in Table S2 using covariates in the Catalist data, which combines public and commercial records of gender, marital status, race/ethnicity, and religious affiliation. In the pooled data, the differences in means are small and generally not significant despite the very large sample size, suggesting the as-if randomization assumption is likely to be satisfied among registrants.¹ As we show in the paper, though, this seeming balance masks consequential differences between the two groups in registration patterns.

¹There are a few imbalances in the 1990 cohort, which may be the result of the shorter interval between the treatment election for this cohort (2008) and the year of data collection (2011).

Table S2: Balance statistics in Catalist data

All	Treatment	Control	<i>p</i> -value
Male	0.464	0.457	0.11
Married	0.095	0.095	0.91
Black	0.160	0.154	0.06
White	0.634	0.644	0.12
Hispanic	0.156	0.159	0.38
Catholic	0.265	0.272	0.06
Protestant	0.275	0.276	0.93
1986	Treatment	Control	<i>p</i> -value
Male	0.461	0.454	0.44
Married	0.126	0.122	0.38
Black	0.156	0.156	0.92
White	0.643	0.645	0.90
Hispanic	0.156	0.156	0.89
Catholic	0.262	0.268	0.35
Protestant	0.283	0.277	0.38
1988	Treatment	Control	<i>p</i> -value
Male	0.472	0.457	0.04
Married	0.085	0.089	0.33
Black	0.161	0.156	0.39
White	0.637	0.645	0.28
Hispanic	0.153	0.153	0.96
Catholic	0.265	0.268	0.65
Protestant	0.274	0.68	0.40
1990	Treatment	Control	<i>p</i> -value
Male	0.461	0.462	0.91
Married	0.070	0.061	0.03
Black	0.164	0.148	0.01
White	0.632	0.644	0.16
Hispanic	0.159	0.174	0.03
Catholic	0.267	0.286	0.02
Protestant	0.269	0.286	0.03


2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792)

S3 RD plots for Study 1


S3.1 Turnout-to-registration rates

As is conventional in RD analyses, we plot raw turnout rates among registrants binned by date of birth. Figure S1 illustrates how turnout varies by eligibility in the election after treatment (which we call E2), while Figure S2 presents corresponding results for the next two elections (E3 and E4).

Figure S1: RD effects of voting eligibility on turnout in subsequent election




(a) Turnout effects in 2006 for 1986 cohort (E2) (b) Turnout effects in 2008 for 1988 cohort (E2)




(c) Turnout effects in 2010 for 1990 cohort (E2)

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792). 95% confidence intervals in parentheses. Lines represent means and 95% confidence intervals for just-eligibles and just-ineligibles.

Figure S2: RD effects of voting eligibility on turnout in second and third subsequent elections



(a) Turnout effects in 2008 for 1986 cohort (E3) (b) Turnout effects in 2010 for 1988 cohort (E3)



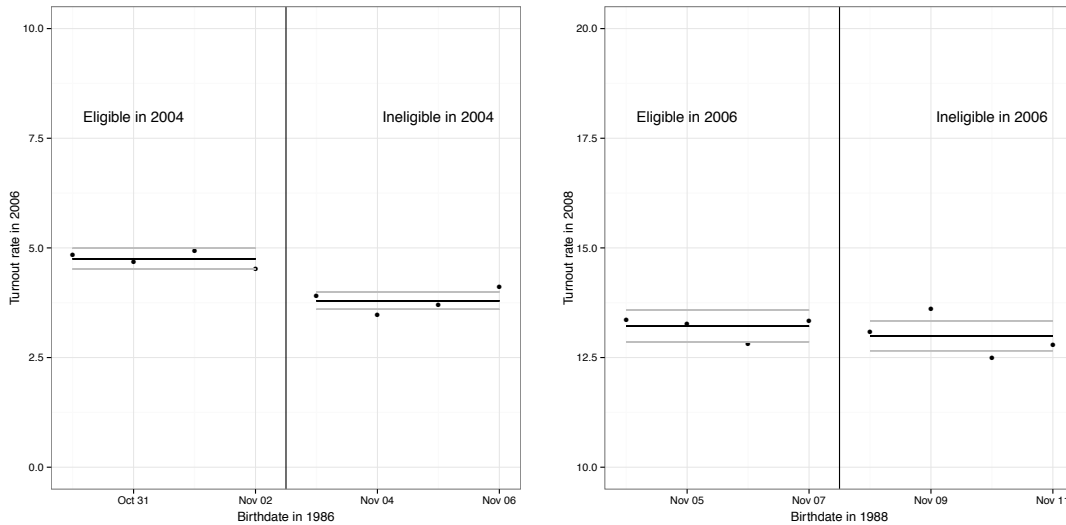
(c) Turnout effects in 2010 for 1986 cohort (E4)

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792). 95% confidence intervals in parentheses. Lines represent means and 95% confidence intervals for just-eligibles and just-ineligibles.


S3.2 Turnout-to-births rates

Figures S3 and S4 plots the raw data for turnout rates by date of birth when adjusted by birth totals rather than the number of registrants in the data.

Figure S3: RD estimates of voting eligibility effects on population turnout rates (E2)




(a) Turnout effects in 2006 for 1986 cohort (E2)(b) Turnout effects in 2008 for 1988 cohort (E2)




(c) Turnout effects in 2010 for 1990 cohort (E2)

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792). 95% confidence intervals in parentheses. Lines represent means and 95% confidence intervals for just-eligibles and just-ineligibles.

Figure S4: RD estimates of voting eligibility effects on population turnout rates (E3–E4)



(a) Turnout effects in 2008 for 1986 cohort (E3) (b) Turnout effects in 2010 for 1988 cohort (E3)



(c) Turnout effects in 2010 for 1986 cohort (E4)

2011 Catalist data; $n = 49,271$ (1986: 18,326; 1988: 17,153; 1990: 13,792). 95% confidence intervals in parentheses.

Lines represent means and 95% confidence intervals for just-eligibles and just-ineligibles.

S4 Results for Study 1 in states without preregistration

Holbein and Hillygus (2015) argue that preregistration increases mobilization for just-ineligible voters, who are exposed to the opportunity to preregister. At the time of the elections that we consider in Study 1, only Florida and Hawaii allowed voters to pre-register when they were 17 years old. Table S3 therefore replicates the main turnout-to-birth analysis of Study 1 excluding Florida and Hawaii from both the turnout and birth counts. The results are largely unchanged from the original analysis.

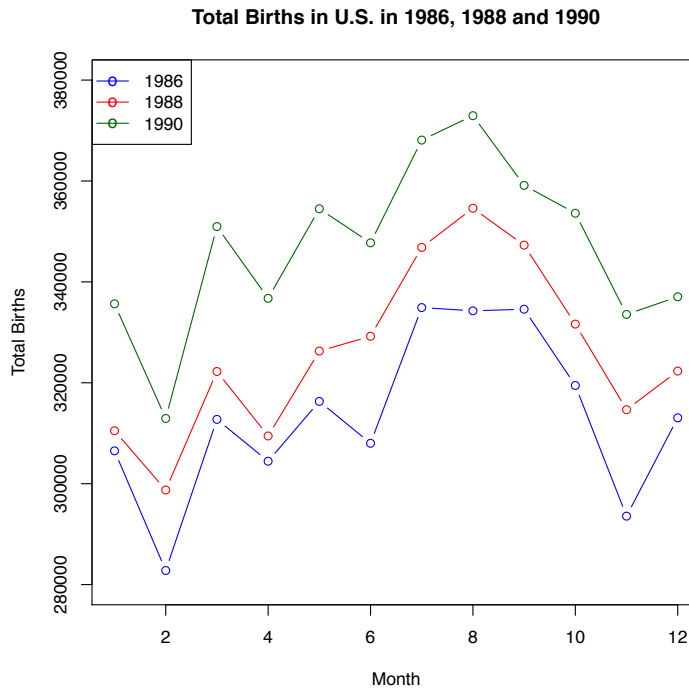
Table S3: Turnout-to-birth rates by voting eligibility excluding preregistration states

A. 1986 cohort (first election for just-eligibles: 2004 presidential)			
	E2 (2006 midterm)	E3 (2008 presidential)	E4 (2010 midterm)
Eligibility effect	0.89 (0.57, 1.21)	0.59 (0.06, 1.12)	-0.28 (-0.62, 0.07)
Control group	3.79	12.93	5.16
B. 1988 cohort (first election for just-eligibles: 2006 midterm)			
	E2 (2008 presidential)	E3 (2010 midterm)	E4 (2012 presidential)
Eligibility effect	0.32 (-0.18, 0.83)	-0.18 (-0.50, 0.13)	-
Control group	12.38	4.54	
C. 1990 cohort (first election for just-eligibles: 2008 presidential)			
	E2 (2010 midterm)	E3 (2012 presidential)	E4 (2014 midterm)
Eligibility effect	1.19 (0.91, 1.49)	-	-
Control group	3.26	-	-

2011 Catalist data; $n = 44,167$ (1986: 16,385; 1988: 15,371; 1990: 12,411). 95% confidence intervals in parentheses.

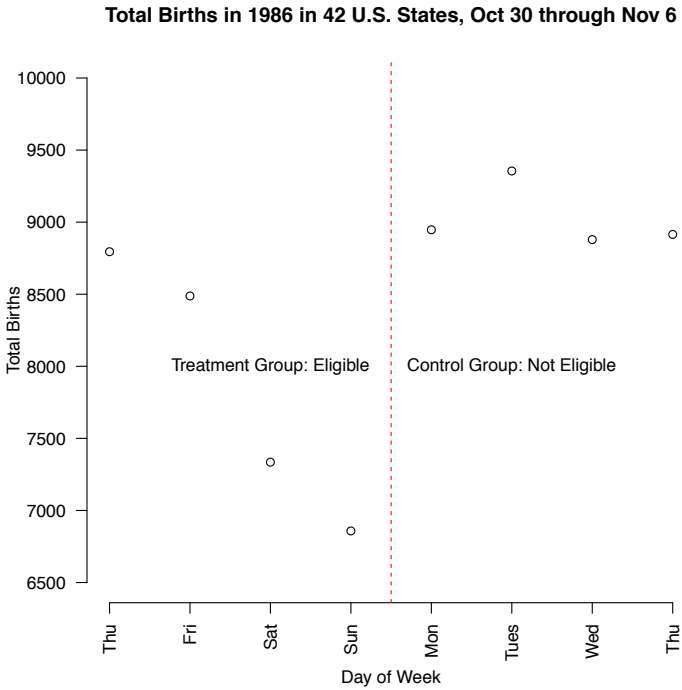
S5 Total births by day and month

Figure S5: Total births by month in the United States

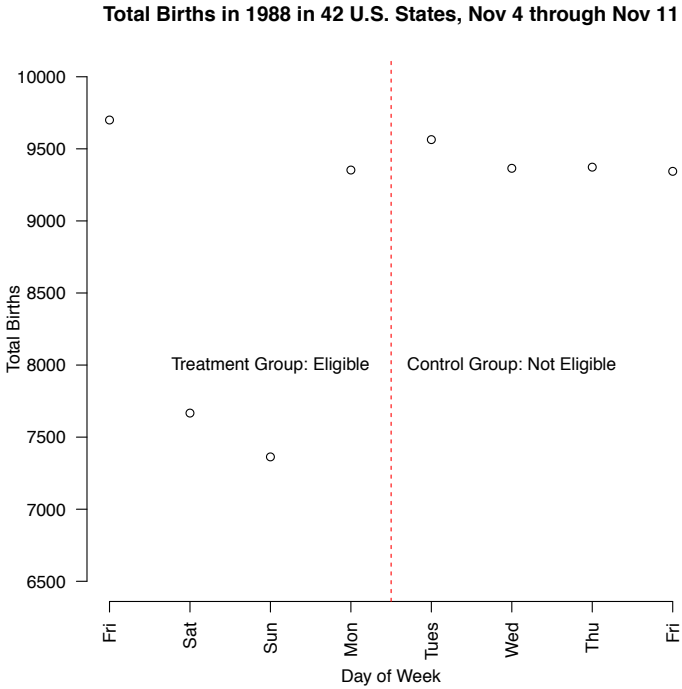


Source: Vital Statistics of the United States for 1986, 1988 and 1990, Volume I, Natality.

Figure S6: Total births by day in 42 U.S. states included in Study 1



(a) Daily birth counts in 1986



(b) Daily birth counts in 1988

Source: Vital Statistics of the United States for 1986 and 1988, Volume I, Natality.

References

Holbein, John B. and D. Sunshine Hillygus. 2015. "Making Young Voters: The Impact of Preregistration on Youth Turnout." *American Journal of Political Science* Early View.