Supporting information for "Trickle-up Political Socialization: the Causal Effect on Turnout of Parenting a Newly Enfranchised Voter"

Jens Olav Dahlgaard*
Department of Business and Politis
Copenhagen Business School

January 7, 2018

*jod.dbp@cbs.dk

Contents

S1 The impact of children coming of age

In the paper, I refer to a previous literature that has proposed children can causally affect parents' turnout and political involvement (Simon and Merrill 1998, McDevitt and Chaffee 2000, 2002, Wong and Tseng 2008). In this section, I postulate that there is still a need for additional causal evidence and review some of the pioneering work on trickle-up political socialization. These studies implemented creative ideas using the best available data. They were path breaking and have inspired the paper in their attempt to call attention to the role that children play in socializing their parents. I provide new and causally identified evidence to support these important claims.

Many previous studies have used the Kids Voting USA program, a civics education program that runs in participating schools prior to elections. Though the program is intended for the participating children, research has focused on its indirect effect on parents, possibly through trickle-up political socialization. Studies have tried to evaluate the effects on parents using primarily survey data and *ad hoc* matching.

Simon and Merrill (1998) found that students and teachers hold positive attitudes toward the program and that turnout is 2.2 percentage points higher in the participating districts. They used a combination of survey data of students and teachers and *ad hoc* matching of districts that participated in the Kids Voting program in 1994 with districts that did not participate. However, the authors themselves suggested that participating areas *"may have had more aggressive social studies administrators whose zeal for such a program energizes civics teaching at all levels"* (Simon and Merrill 1998, 39). Similarly, one cannot rule out that their matching technique fails to match on all other relevant observed and unobserved predictors of turnout.

McDevitt and Chaffee (2000) found that parents whose children participated in the program in 1994 in San José, California, had higher political knowledge. They relied on the assumption that participation in the program is independent of children's and parents' characteristics since teachers were the ones who chose for their classroom to participate. However, one could imagine how teachers opting in could be teaching in areas where for one reason or the other parents were more likely to vote regardless of the program. Furthermore, they use intensity of treatment exposure as a continuous independent variable and measure it by self-reported recall, which could be correlated with other characteristics such as political interest and knowledge. Therefore, one is not fully assured that their independent variable, recall of exposure, is truly randomly assigned, which in turn may limit the causal conclusions of their study.

Two other studies emerging from the Kids Voting program are Linimon and Joslyn (2002) and McDevitt and Kiousis (2006). Linimon and Joslyn (2002) used aggregated turnout across counties in Kansas to show that counties adopting the Kids Voting program exhibited higher turnout in the 1996 presidential election than counties that did not adopt the program. McDevitt and Kiousis (2006) use panel survey data of children and parents over several years and focus groups to track the long-term impact of the program in four counties across the US.[S1] However, both studies are subject to variations of the same issues as listed above, such as an inability to rule out self-selection by adopting counties and self-reported measurement of exposure to the program.

Trickle-up socialization has also been studied outside the Kids Voting program. Using two rounds of survey data from adolescents in Lubbock, Texas, surveyed before and after the US 2000 Presidential Election, McDevitt (2006) finds a correlation between the adolescents' self-reported degree of initiating political discussions at home and the feedback in the form of political stimuli they report receiving from their parents, which could influence both adolescents' and parents' political attitudes. However, this study too relies on self-reports of both independent and dependent variables and there is no clear exogenous shock to the childrens' propensity to initiate discussions.

Turning the attention to socialization processes in immigrant families, Wong and Tseng (2008) find support for their model of bi-directional socialization in their study that mixes survey data and in-depth interviews. Relying on qualitative interview data, Bloemraad and Trost (2008) argue that for immigrant families political socialization happens in the reverse with parents being influenced by their offspring. However, though both articles make careful use of their in-depth data, their data is eventually dependent on self-reports and self-evaluation, subject to the same potential biases as when used in a survey setting.

The review demonstrates that though there are both theoretical arguments for trickle-up socialization and empirics that do not contradict it, a skeptic could object that there is no clear causal appraisal of the theories regarding the coming of age of adolescents and parents' political behavior. Besides the issues pointed out above, another point is that the studies are overwhelmingly from very local contexts, many of which opted in to some treatment such as the Kids Voting program. Since there is self-selection in participation, the conclusions are sensitive to potential bias. For instance, one could imagine how officials choose to participate if they believe that the program will be beneficiary to their context and opt out if they believe it to be harmful. The analyses in the paper provides causal appraisal for these important claims.

[S1] Another well-cited article is McDevitt and Chaffee (2002). However, that article primarily presents a model for trickle-up socialization and only hint to empirical findings from the authors' past studies.

S2 Means of characteristics for voters in municipalities that participated in different years

In the paper, I rely on data from all municipalities only in 2013, while I have to rely on data from a subset of the municipalities in 2009, 2014, and 2015. I present the effects using data from all municipalities with available data as well as for only municipalities for which I have data in all four election. As I point out in the paper, the local effects are causally identified in all four years for the participating municipalities, but the municipalities for which I estimate the effect vary from election to election, and one can question what that means for the conclusions. In this section, I turn the attention to descriptive statistics for the municipalities that participated each year, and in all four years, to see if we can find clear indications that the municipalities varies on other characteristics than their decision to participate or not. In section S6, I go into more detail with respect to which municipalities I have data on for each year.

In Table S1, I show the means of interesting variables as of 2013 for different subpopulations. The first column includes parents with children turning 18 within a window of six months around the age threshold for eligibility in all muncipalities that participated in 2013. The second column has descriptives for all voters in municipalities that participated in 2009, the third column for those participating in 2014, and the fourth for those participating in 2015. Finally, the last column shows descriptives only for those who lived in municipalities that participated in all four election studies. I emphasize that all columns show the descriptive statistics in 2013 based on living in a municipality that was in the data in a given year. I chose this because we can better make the comparison when there are no time trends to take into account. I also emphasize that the groups are not mutually exclusive. For instance, the voters that live in the municipalities participating in all four years are also in the summary for each individual year.

Importantly, Table S1 reveals only small differences in the composition of voters. If we only compare the first and the last columns, we see that voters in the consistently participating municipalities tend to have lower turnout, are less likely to be native Danes, and more likely to have at least a bachelor's degree. These differences are largely driven by the fact that some of the large urban municipalities, Copenhagen being by far the biggest, are included in all of the elections, and they have lower turnout, more diversity, and better educated citizens.

Table S1: Means for parents in 2013 based on living in participating municipalities over election years

	2013 (All)	2009	2014	2015	2009-2015
Parent voted (%)	77.25	75.95	76.70	76.73	75.34
Parent is native (%)	88.07	85.69	87.38	86.96	84.94
Parent's age (%)	44.59	44.67	44.63	44.64	44.72
Parent's income (1000 euro)	57,876	57,608	58,804	58,531	58,190
Parent is female (%)	51.08	51.03	51.09	51.12	51.07
Parent has bachelor's degree (%)	27.02	28.39	27.37	27.78	28.90
Child is female (%)	48.22	47.94	48.37	48.24	48.05
Child is oldest (%)	44.60	45.02	44.93	44.73	45.19
N	132,003	70,310	85,396	104,490	51,313
No. of Municipals	98	44	61	72	32

For each variable I take the mean for voters in 2013 who lived in a municipality that was included in the election in each election year.

S3 Robustness of RD findings

There are several ways to assess the robustness of the RD design. In this section, I use a wide range of alternative estimation strategies, none of which change the conclusion that parents are mobilized to vote when their children reach voting age. First, in the left panel of Figure S1, the model is estimated at all bandwidths from three to 183 days.[S2] I estimate the effect independently in each year using the chosen bandwidth instead of letting the algorithm pick the bandwidth. I take precision weighted averages of the effects. The pooled estimates range from approximately 1.9 to 4.8. In addition, the confidence intervals only overlaps zero in the smallest bandwidths where the estimates are very imprecise.

Another test is to estimate the effect at alternative cutoffs, where we would expect to see no jump. In the right panel of Figure S1, the effect is estimated using all monthly cutoffs for the age of the child from 14 years to 22 years. Only in the true cutoff are we comparing parents of all eligible children to parents of all ineligible children. The effect estimates are dispersed around zero and except for one outlying estimatate, which is large and in the opposite direction, the numerical values of the sizes of the mock effects are far from the true effect, which is indicated by the dashed, horizontal lines.

That the estimate in the true cutoff stands out finds additional confirmation in Figure S3. The plot extends the robustness test presented in Figure S2. Here, the mock estimates are compiled as a histogram in the first panel, with the dashed line being the effect in the true cutoff and the negative value of it. In the second panel, numeric t-values are plotted and the dashed line is the t-value in the true cutoff. Here, it is still only the one outlying, large estimate that compares in size, although not direction, with the effect in the true cutoff.

[S2] In smaller bandwidths the model is not statistically identified.

Figure S1: The effect of child eligibility with different bandwidths

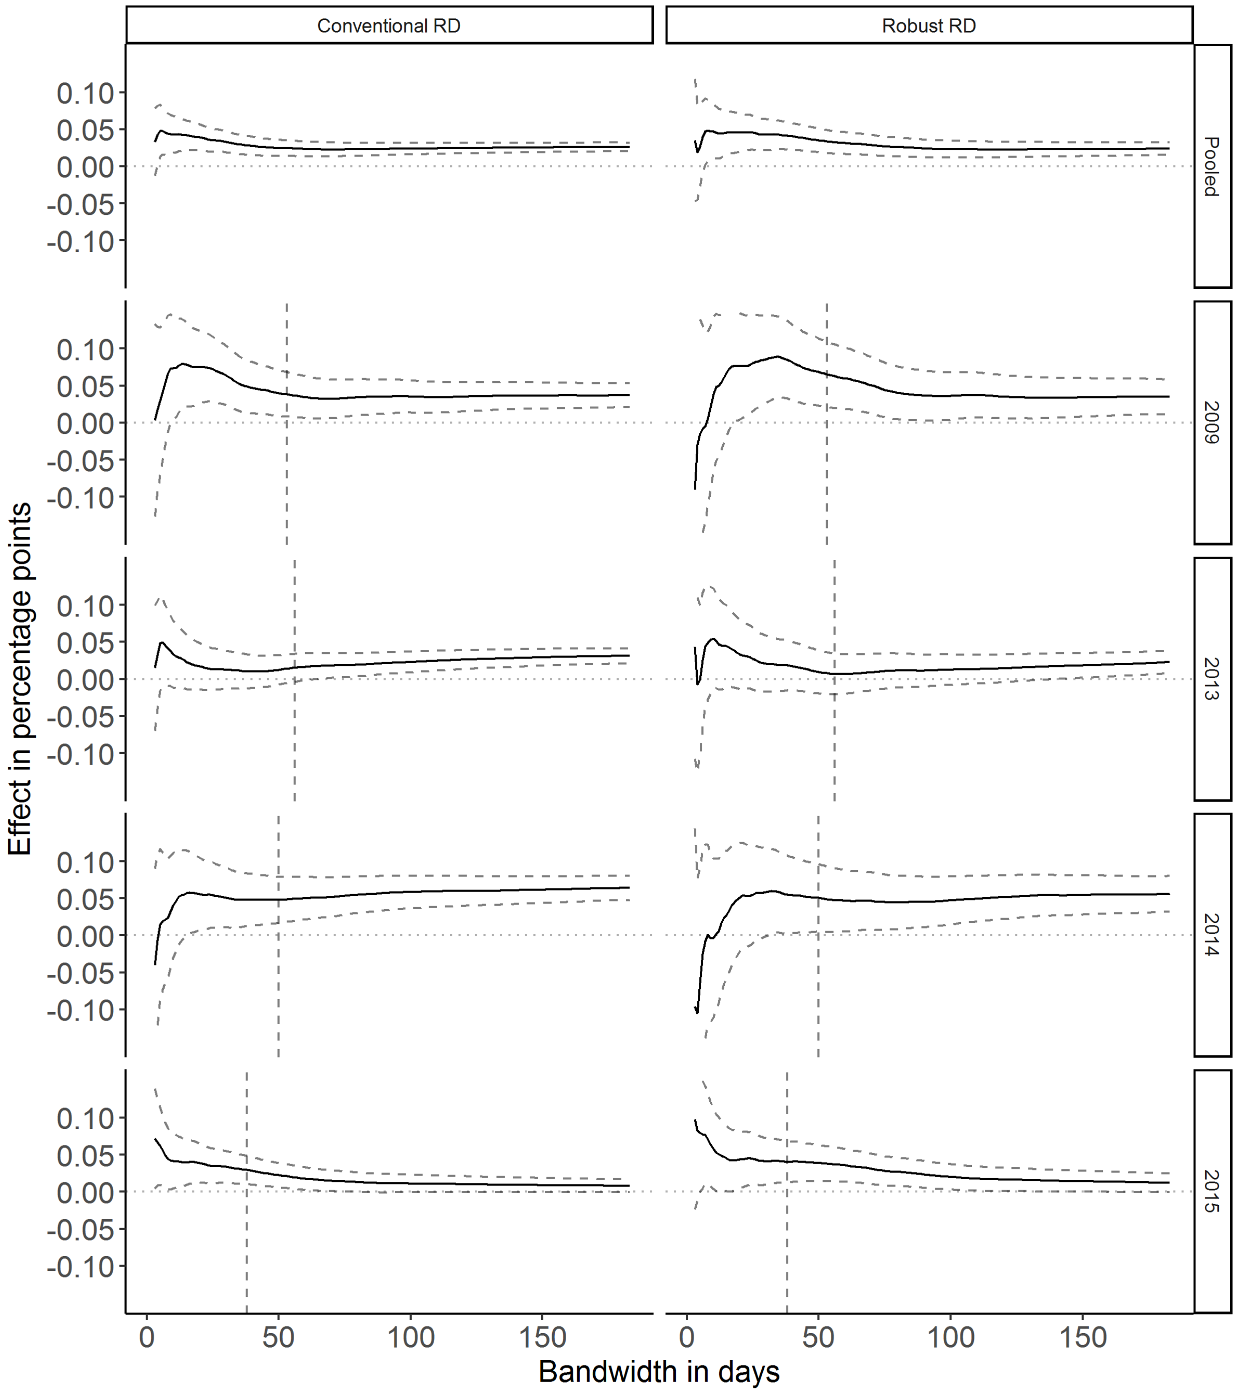

The figure contains model estimates and 95 percent confidence intervals at all bandwidths from three to 183 days. In smaller bandwidths, the **rdrobust**-package cannot estimate the effect.

Figure S2: The effect of child eligibility with different cutoffs

The figure contains estimates in all monthly cutoffs of child age on Election Day from 14 to 22 years, including the true cutoff, 18 years. The dotted lines represent zero effects. The dashed lines in the right panels are the positive and negative values of the effect in the true cutoff, and the bar around the effect in the true cutoff shows the 95 percent confidence interval.

Figure S3: Distribution of estimates at alternative cutoffs

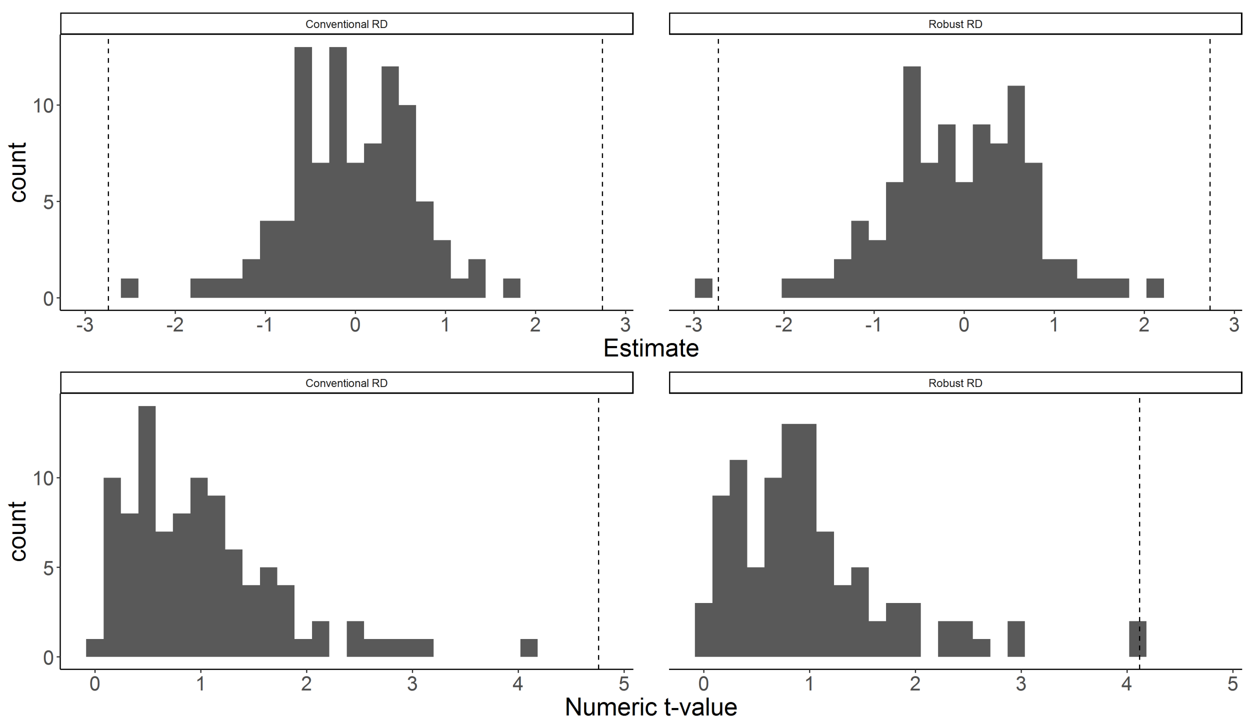

The dashed lines are the effect estimate around the true cutoff and its negative value in the top panel and the t-statistic in the bottom panel.

S3.1 Heterogeneous effects conditional on cohabitation in municipalities with data in all four elections

For the municipalities with data in all four elections, we can also replicate the relationship between the effect size and cohabitation. In figure S4, I plot the relationship between turnout among untreated parents and the effect size conditional on cohabitation with the child. The plot is similar to the one I presented in Figure 2 of the paper, but here I only include voters with data in all four elections. The figure reveals the exact same pattern as when I include all municipalities in the main paper. The confidence intervals around the point estimates are wider, since we lose data points from all the municipalities that do not have data in all four elections.

Figure S4: Heterogeneous effects conditional on cohabitation and level of turnout in municipalities with data in all four elections

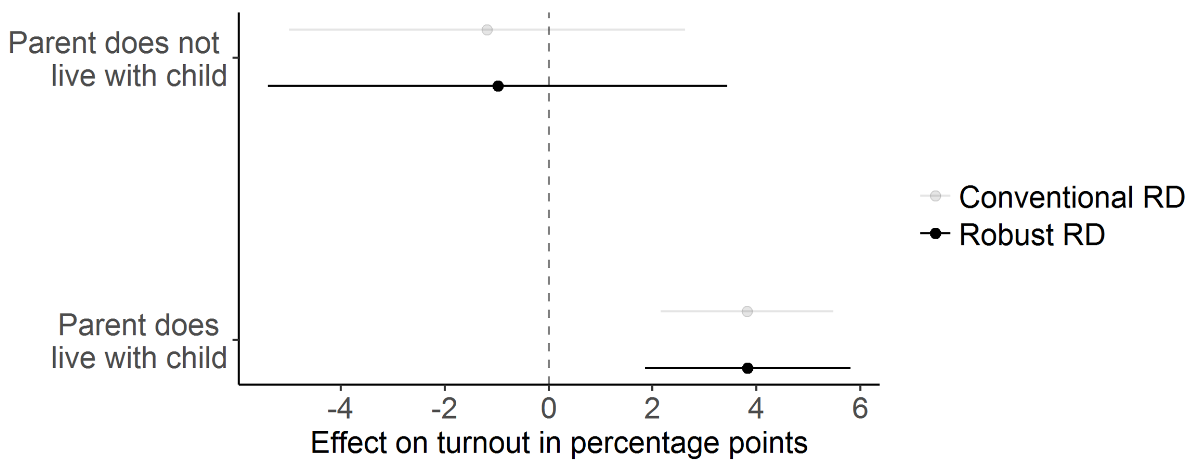

The figures contain point estimates and 95 percent confidence intervals for the effect for parents that live with their child (gray) and parents that do not (black) across all four elections. The model used to estimate the heterogeneous effects differs from the model in the main analysis in that it use data aggregated per week. The results are robust to using daily estimates on data that cannot be made available for replication due to privacy restrictions.

S3.2 Controlling for individual covariates and clustering by child

One concern by using the data aggregated by day is that most children will have two parents, both of whom can vote. We could worry that we underestimate the standard errors when we fail to account for that. In this section, I show that the results are robust to clustering standard errors by child. The results I present in this section rely on data that will not be made available for replication. Due to privacy concerns, Statistics Denmark has restrictions on what kind of data can be shared. As we will see the models, where we cluster by child give the same conclusions as the models, where we only identify parents by their children's date of birth, which means that the priority of making replication data available does not come at the cost of overly confident results.

Table S2 contains two models clustered by child. The first model is the same as the main model, except for how I cluster the standard errors. The second model is expanded with a battery of prognostic, background covariates: the parent's age, the parent's age squared, education, income, ethnicity, sex, turnout in the latest election, turnout for the partner of the parent in the past election, municipality fixed effects, the child's sex, if the child is the first-born of the parent, and the child's ethnicity.[S3] In all models, parents are clustered by the identity of their child. As for the analyses in the paper, the models are estimated for each year and pooled together by taking precision weighted averages.

Under the assumptions of the regression discontinuity design, the local effect is causally identified. Therefore, the estimate without controls should be unbiased and including controls should not have a large impact on the point estimate. The advantage of including pretreatment covariates is that if they are predictive of the outcome, they can reduce variability and lead to a more precisely estimated point estimate (Calonico et al. 2017, 375). It is also arguably the case that if the point estimate changes dramatically, it could indicate that the assumptions are not met and that the design is invalid.

[S3]Means are imputed for missing values and a dummy for being missing is included.

Table S2: Individual level models with standard errors clustered by child

	No control for covariates					Control for covariates				
	2009	2013	2014	2015	Pooled Effect	2009	2013	2014	2015	Pooled Effect
Estimate	3.62	1.64	4.89	2.70	2.71	4.39	1.24	5.11	2.60	2.75
St.err.	1.74	1.08	1.84	1.04	0.64	1.64	1.06	1.63	0.98	0.61
CI	[0.21, 7.03]	[-0.48, 3.76]	[1.28, 8.50]	[0.66, 4.74]	[1.46, 3.96]	[1.18, 7.60]	[-0.84, 3.32]	[1.92, 8.30]	[0.68, 4.52]	[1.55, 3.95]
Robust estimate	3.75	1.28	4.47	3.11	2.70	4.63	0.95	4.83	2.99	2.79
Robust st.err.	2.09	1.25	2.18	1.18	0.74	1.96	1.25	1.94	1.11	0.71
Robust CI	[-0.35, 7.85]	[-1.17, 3.73]	[0.20, 8.74]	[0.80, 5.42]	[1.25, 4.15]	[0.79, 8.47]	[-1.50, 3.40]	[1.03, 8.63]	[0.81, 5.17]	[1.40, 4.18]
$\text{Turnout}_{parents\ of\ ineligible}$	68.15	75.34	58.38	88.93	74.48	68.20	75.23	58.45	88.98	74.48
$N_i neligible$	9,285	19,961	10,044	10,910	50,200	9,285	17,950	11,193	10,383	48,811
$N_e ligible$	9,368	21,396	9,746	10,455	50,965	9,368	19,115	10,852	9,930	49,265
Bandwidth	57	60	54	42		57	54	60	40	
Bias-correction bandwidth	86	100	85	76		87	87	92	74	
Individual level control variables	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes

The first model is similar to the model I estimate in the paper, only parents are clustered by their child. In the second model, I include a range of prognostic, pretreatment covariates. The models use individual level data that due to privacy concerns cannot be made publicly available. The pooled values are precision weighted averages.

The models give estimates that are similar to each other and to the estimate reported in the paper. In model 1, the pooled, conventional point estimate is 2.71, instead of 2.80 as reported in the paper. The 95 percent CI is [1.46, 3.96] compared with [1.66, 3.94] in the paper. The small deviation stems from the fact that the models are estimated for each year before they are pooled together, and use individual level data. The precision weighted average assigns to each estimate a weight proportional to its precision, which is the reciprocal of the variance. When I cluster the standard errors by the child instead of by the day, the variance of each estimate slightly changes, which in turn leads it to being assigned a different weight. The individual level data with clusters also means that the algorithm picks slightly different bandwidths, which in turn causes small changes to the point estimates.

We also see that the standard error is slightly larger when I cluster by day. This is also as we would expect. Instead of 0.58, as I report in the paper, the standard error when I cluster by child is 0.64, which indicates that in the paper, I overestimate the precision of the estimates slightly. Substantially, the results are the same when I use the data with errors clustered by week. I favor the fact that I can make this data available for replication over the underestimation of standard errors.

In model 2, I include prognostic background covariates. When we include background covariates, we would expect to see the estimate change only by sampling variation, and to see slightly smaller standard errors due to a smaller regression error. This is consistent with what we see. The estimate changes to 2.75 with a 95 percent CI of [1.55, 3.95]. In conclusion, the substantial results do not change when running the model on individual level data where we can cluster by day, and the divergence when we include a battery of prognostic pretreatment covariates could just be due to sampling variation.

S3.3 Alternative polynomial specifications

In the paper, I use a linear fit of turnout on the children's age in days, which is the default option in the software I use. In Figure S5 I plot the effect in the cutoff if I specify the fit as anything from a first order polynomial (as in the paper) to a sixth order polynomial, although the literature advises against fitting higher order polynomials for RD designs (Gelman and Imbens 2017, Gelman and Zelizer 2015). Once again, we conclude that under any alternative specification, the sign of the effect is the same and the confidence intervals never comes close to including zero. In fact, the point estimate is larger for each alternative choice for the order of the polynomial.

Figure S5: Pooled effect with different polynomial fits

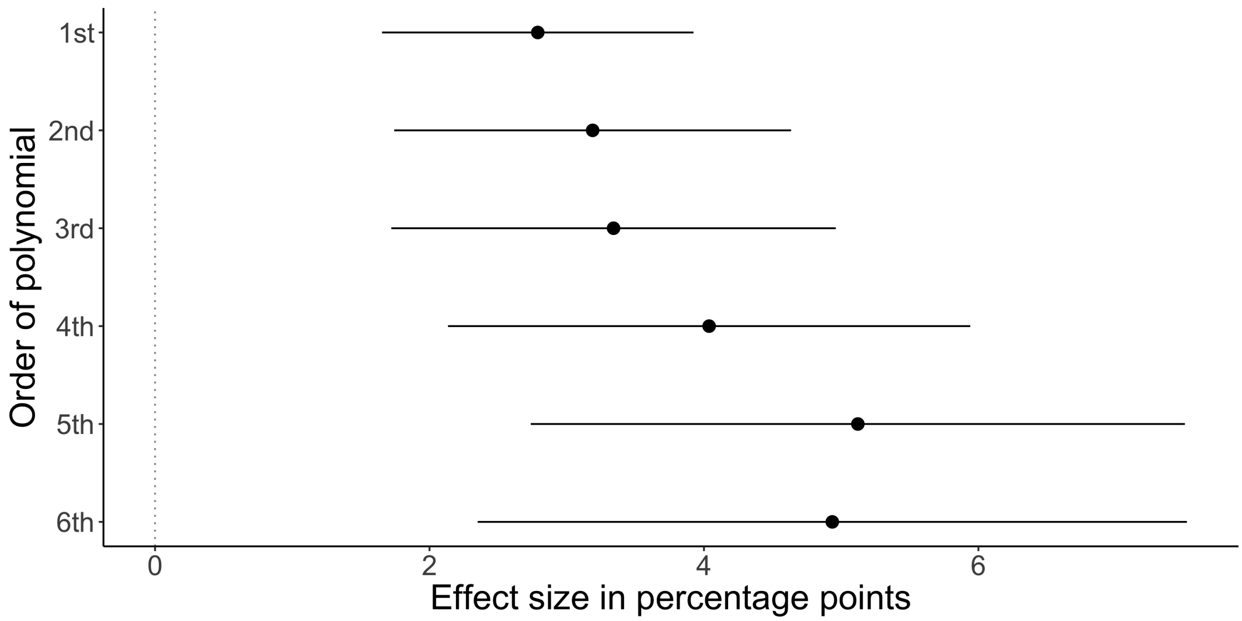

The first model estimate is the same as I estimate in the paper. All the other estimates are similar only the order of the polynomial fit is increased.

S3.4 Estimating the effect on other variables

As another test of the robustness of the RD design, I substitute the dependent variable with the dichotomous background covariates from Table 1 in the paper. When substituting turnout with other variables, we would expect the jumps in the cutoff to be small and the confidence intervals to overlap zero. In Figure S6, I plot both conventional and robust RD estimates for each year and pooled for each of the six covariates. The figures reveal that for all sixty estimates but one, the confidence interval includes zero. The only exemption is the conventional estimate for whether the child is the oldest of the parent's children in 2014. Considering that we have sixty estimates on what we assume to be variables unaffected by the treatment, I see it as evidence for the claim that only one confidence interval excludes zero.

Figure S6: Effect of child eligibility on other variables assuming continuity around the cutoff

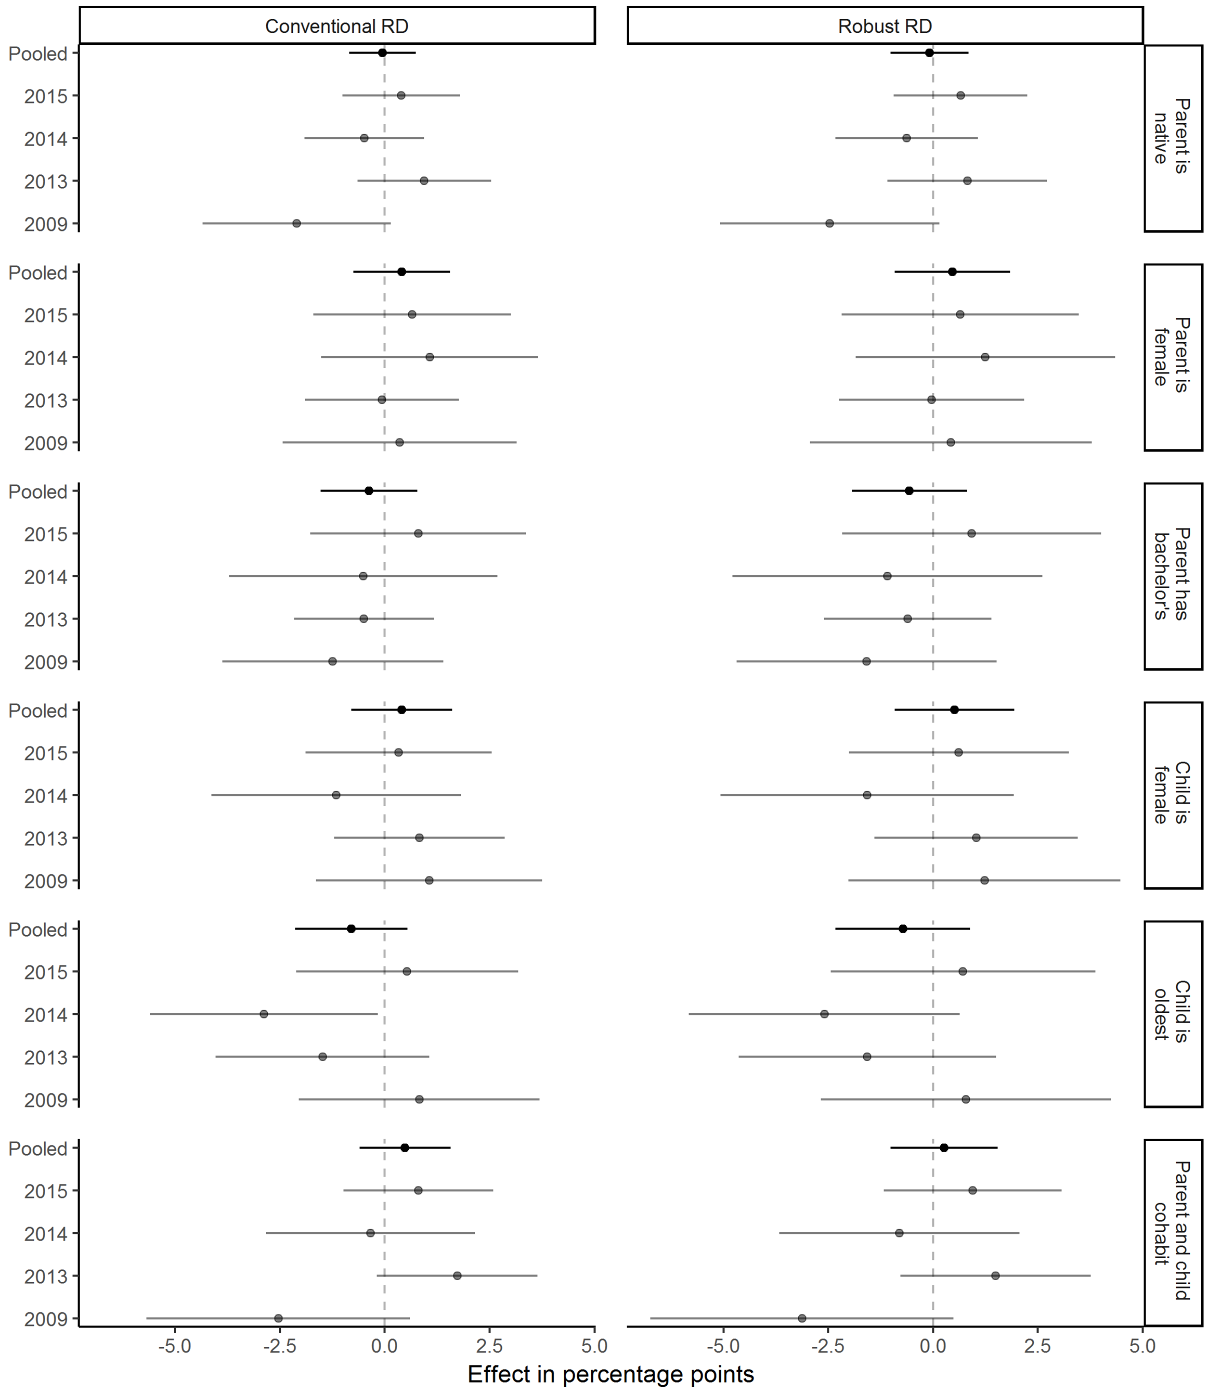

NOTE: The estimates are precision weighted averages. The confidence intervals are 95 percent confidence intervals

S3.5 Estimating the effect under the assumption of local randomization

In the paper, I also analyze the data as a local experiment in windows of one day, two days, three days, one week, and a month. In Figure S7, I show the effect estimates and confidence intervals for parents of 18 year-olds in all estimation windows up to 183 days (six months). As mentioned in the paper, we may underestimate the confidence intervals, because I am basing the analyses on aggregated data that can be made publicly available. Therefore, I use both 95% and the more conservative 99% intervals.

Based on the figure, every reader is free to choose whatever window he or she finds most plausible. Regardless of what window this is, Figure S7 show that the effect is also robust to being analyzed as an experiment, as we see a positive effect estimate in even the one day window, and the confidence interval in the two day window already excludes zero.

Figure S7: Difference in the proportion of voters over bandwidth in days (averaged over all elections)

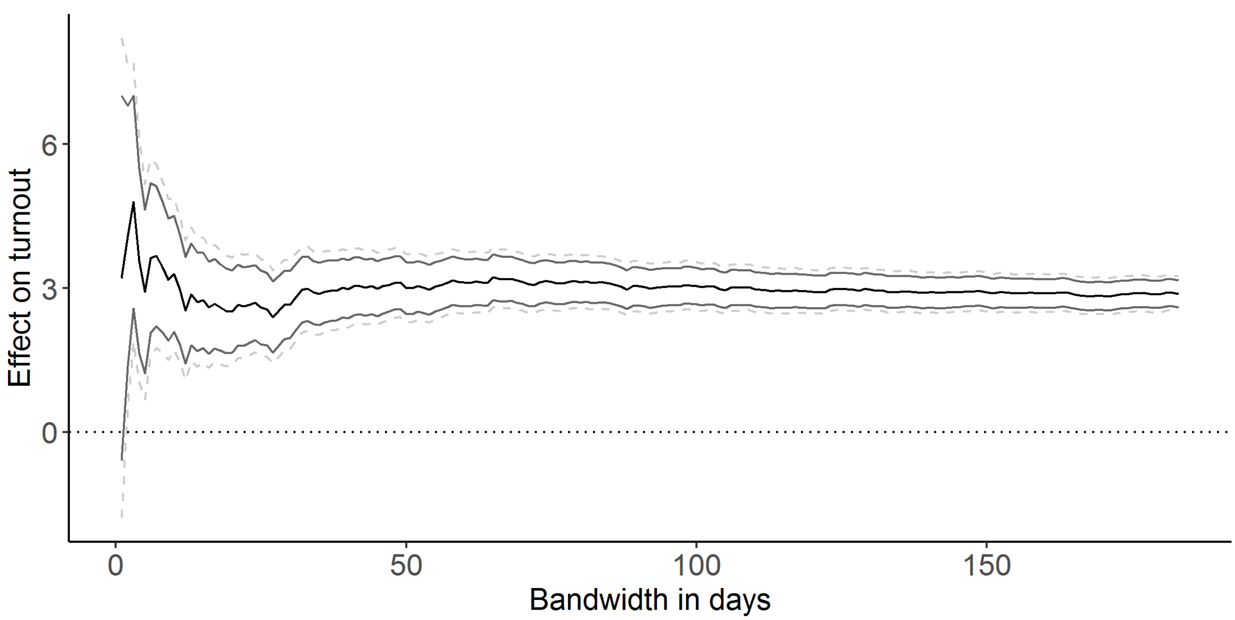

The black line shows the estimated difference in proportion between parents with eligible and ineligible children with 95% CI (gray) and 99% CI (gray and dashed) over bandwidth in days.

To further substantiate the claim that comparisons in the small bandwidth are valid, in Figure S8 and Figure S9 I show p-values for differences in proportions between parents above and below the cutoff on six other covariates. These are if the parent is a native, female, has a bachelor's degree, lives with the child, if the child is the oldest child, and if the child is female. I show the p-values for the balance on these covariates in windows of 183 days in Figure S8, and in Figure S9 I zoom in on ten days windows. We see that only the p-value for whether the child is oldest drops below 0.15 in the ten day window, but it remains above 0.05.

In Figure S10, I expand on the analyses in Figure S8 and Figure S9 as I estimate the difference in proportions for all years and all six pretreatment covariates in windows of one day, two days, three days, one week, and a month. Because the precision of the estimates increases markedly with the estimation window, I allow the x-axis to vary for the different facets of the plot. There are no clear imbalances that stand out for any of the covariates in any of the estimation windows. In line with the previous robustness checks, I see this as supporting the validity of the RD design. Specifically, it supports an as-if random interpretation of the RD design.

I also plot the effect on parents' turnout in each year and pooled for all elections. The pooled estimate corresponds to the effect that I present in the paper's Table 3. As in the paper, we see positive, pooled point estimates in all windows. The yearly estimates are also positive, but of varying size in the all windows, except for the estimate from the European Parliament election of 2014. The point estimate in the one day window for that year is negative, but also very imprecisely estimated.

Finally, in Table S3, I expand Table 3 of the paper by also including a confidence interval on the distribution of parents above and below the cutoff. I simply estimate the proportion of parents that have eligible children and the confidence interval around the estimate. We see that it is only in the widest window of one month, that 50 (which indicates an equal distributions of parents) is not included in the confidence interval.

Table S3: Difference in the proportion of voters over bandwidth in days (averaged over all elections)

Bandwidth	1 day	2 days	3 days	1 week	1 month
Effect on turnout	3.21	4.09	4.8	3.67	2.66
95% CI	[-0.60, 7.02]	[1.37, 6.81]	[2.58, 7.02]	[2.22, 5.12]	[1.96, 3.36]
$\text{Turnout}_{parent\ of\ ineligible}$	74.77	74.09	73.32	73.62	74.54
$\text{N}_{parent\ of\ ineligible}$	979	1,922	2,882	6,707	28,813
$\text{N}_{parent\ of\ eligible}$	931	1,847	2,848	6,724	28,204
$\text{CI}_{distribution of N}$	[46.50, 50.99]	[47.41, 50.60]	[48.41, 51.00]	[49.22, 50.91]	[49.06, 49.88]

Figure S8: Two-sided p-values of imbalances on other variables with bandwidths from one to 183 days

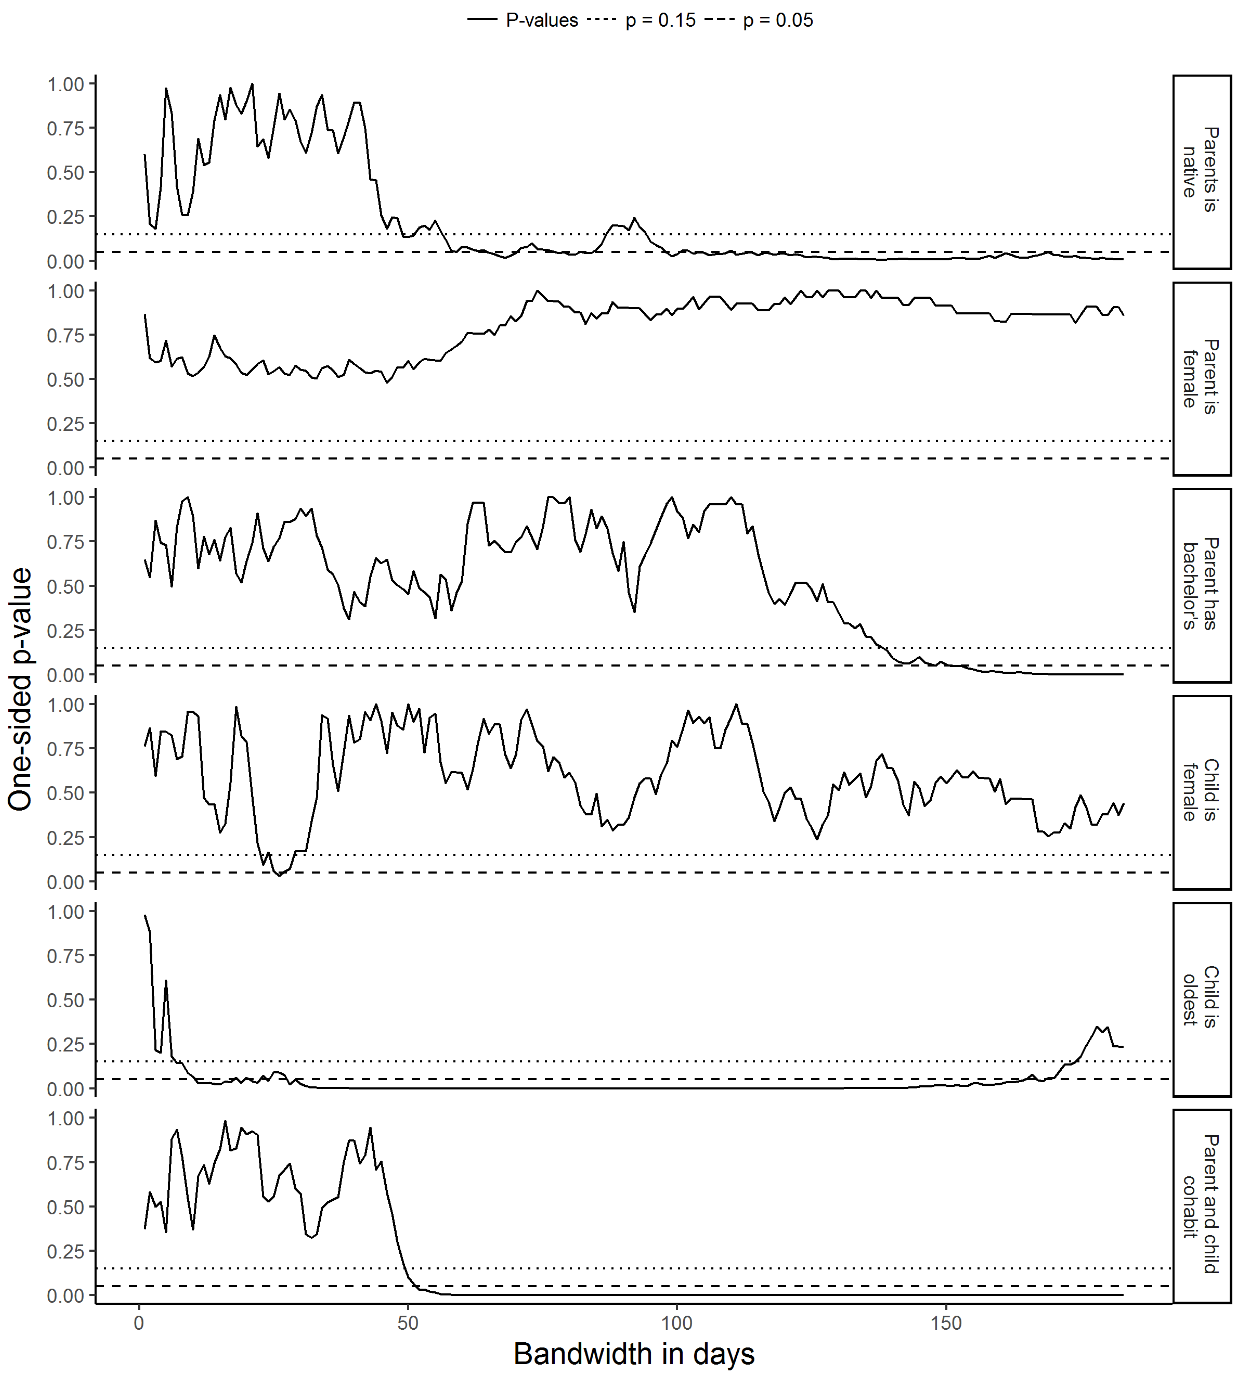

NOTE: The p-values are from a two-sided test of differences in proportions with data pooled for all years.

Figure S9: Two-sided p-values of imbalances on other variables with bandwidths from one to 10 days

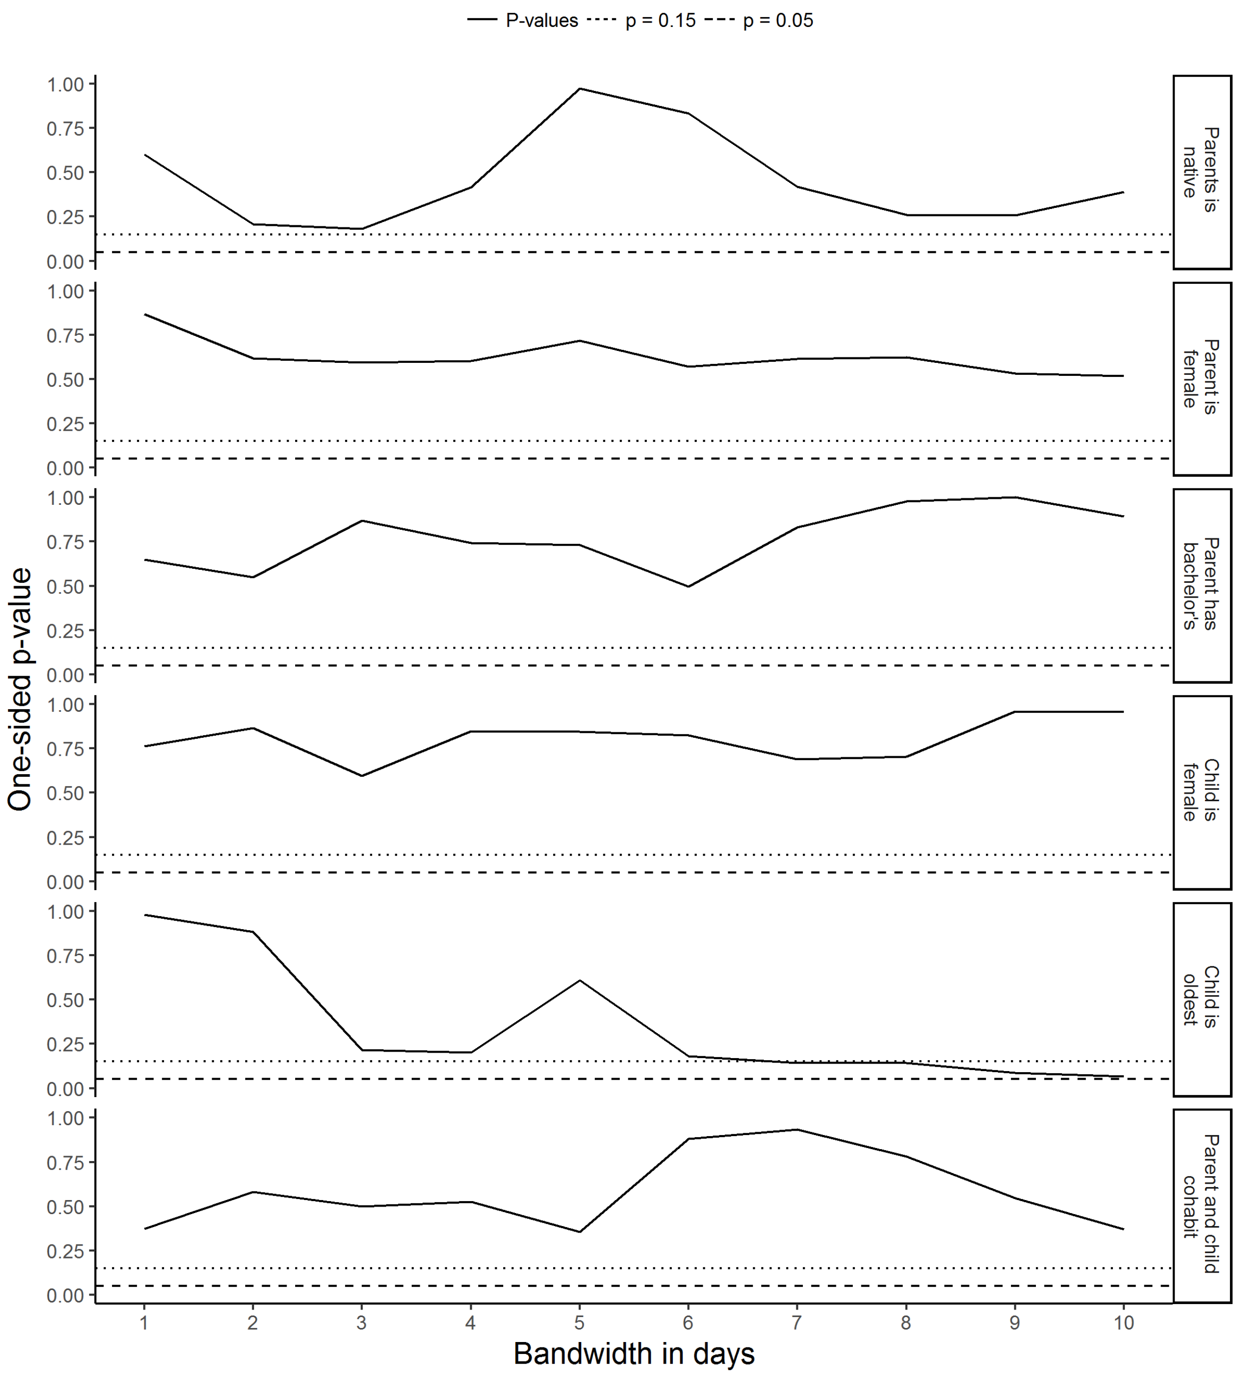

NOTE: The p-values are from a two-sided test of differences in proportions with data pooled for all years.

Figure S10: Effect of child eligibility on other variables assuming random assignment around the cutoff

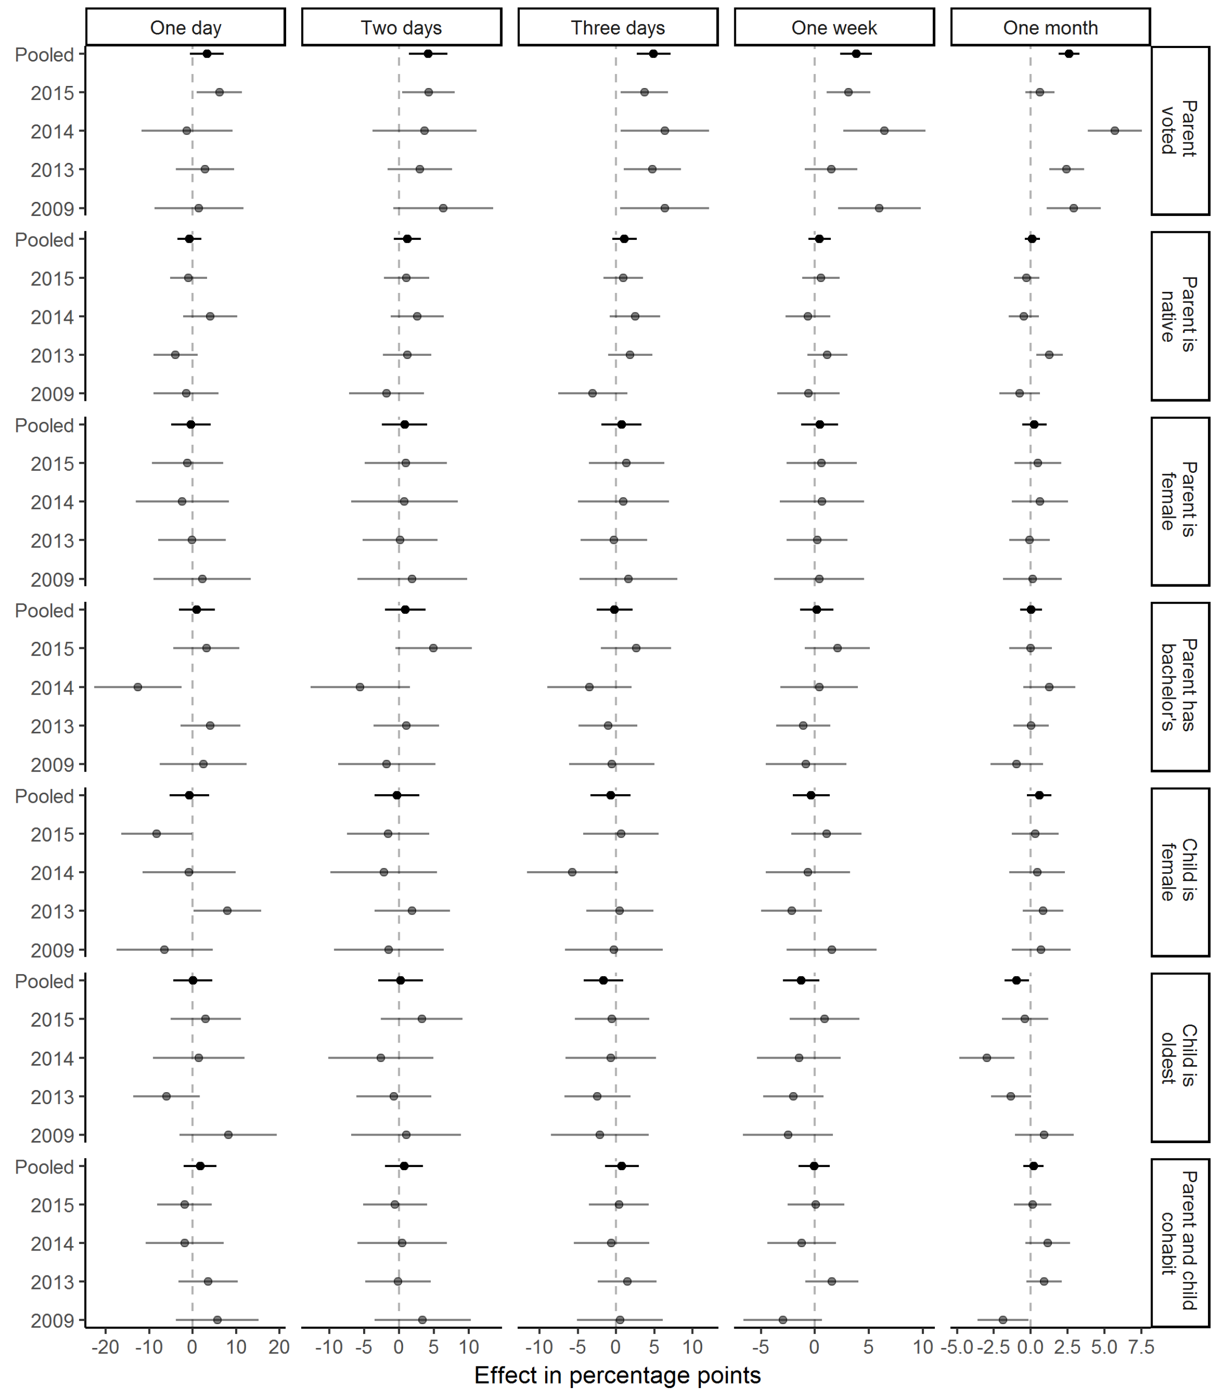

NOTE: The estimates are from a two-sided test of differences in proportions with data pooled for all years. The confidence intervals are 95 percent confidence intervals

S4 Heterogeneous effects and potential mechanisms

As I discuss in the paper at least two mechanisms can explain the effect. Parents whose child has recently become eligible to vote might be keen on inculcating the norm of voting with their child and thus be motivated to turn out on Election Day. Alternatively, children might independently of their parents' inclinations be excited about the prospect of voting, and their excitement spills over to their parents, which I called the trickle-up mechanism. Here, we look at heterogeneous effects to learn more about the extent to which either mechanism drives the effect.

In the paper, I look at how effects vary conditioned on cohabitation, and weI summarize some additional findings based on past voting and whether the parent has a high school diploma and if the child is in high school. In this section, I show these findings. In Figure S11, I split the effects conditioned on turning out in earlier elections. In the first split, I condition on voter turnout in the 2009 municipality election when estimating the effect in the 2013 municipality election. In the second split, the effect in the 2014 European election is estimated conditional on voting in the 2013 municipality election.[S4] These splits create groups of likely and unlikely voters. In 2014, turnout was as low as 20 percent among those who did not vote in 2013 while it was close to 70 percent for those who did vote. In 2013, turnout was just above 45 percent for those who abstained in 2009 while it was 89 percent for those who voted (see Table S4).

If the effect is primarily driven by parents who have internalized a strong norm of voting and want to pass that on to their children, one could expect that effects are stronger among parents who voted in the past elections. After all, these are parents that have previously complied with the norm of voting in an election of either similar or higher salience. On the other hand, if young voters pushing their parents to the polls primarily drive the effect, we might expect to see the strongest effects among the parents who, by previously abstaining, have indicated that they have not internalized the norm of voting.

The results from Figure S11 suggest that both mechanisms are at play, since effects are similar irrespective of past behavior. In the 2014 European election, the effects are comparable for parents regardless of their behavior in 2013, both when the effects are expressed in percentage points and probits. In 2013, we see

[S4] A full set of possible spillovers would include the effect in 2014 conditioned on voting in 2009 and the effect in 2015 conditioned on voting in 2009, 2013, and 2015. However, the included comparisons are chosen because they compare effects based on the most recent election with available data. Second, these splits give effects in an election with similar turnout to the one conditioned on (2009 to 2013) and where turnout is higher in the election conditioned on (2013 to 2014), which makes it possible to create groups of arguably high and low social norm voters and examine their behavior in a lower and equal salience election.

Figure S11: Assessing heterogeneous effects for previous abstainers and voters

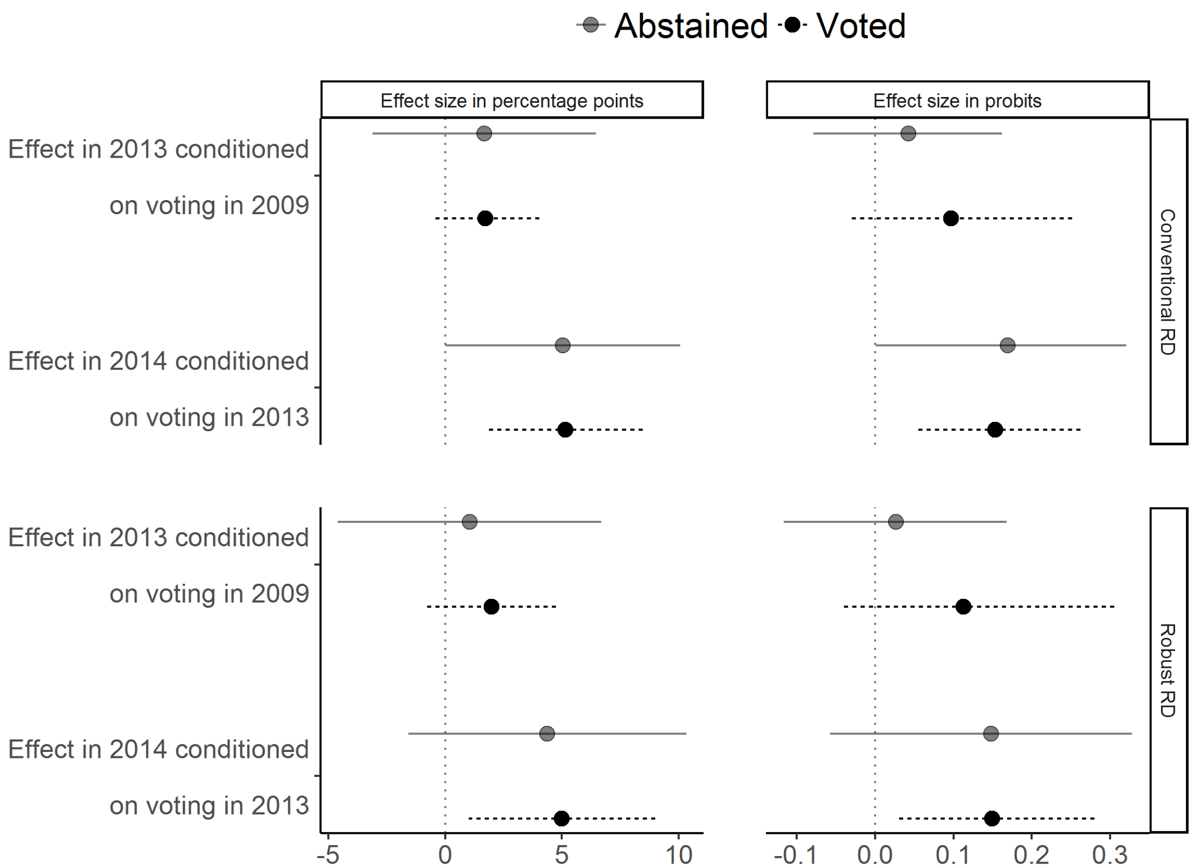

The model used to estimate the heterogeneous effects differs from the model in the main analysis in that it uses data aggregated per week. The week aggregation is chosen to assure individual anonymity in the replication data. The results are robust to using daily estimates (not reported).

similar effects between voters and abstainers from 2009. In conclusion, both parents who usually vote and parents who usually abstain are driving the aggregate effect. If the first mechanism drives the effect, we would expect the first type to be mobilized. If the second mechanism drives the effect, we would expect the second type to be mobilized. When both types increase participation, it suggests that both mechanisms are at play. In Table S4, I show baseline turnout rates and group sizes for the parents included in Figure S11.

As a another attempt to pin down the mechanism, we can look at the effect conditional on the parents having high school or more as highest educational level and the children being in high school. Previous research has demonstrated that young people in high school are more likely to vote than their peers who are not in high school (Bhatti and Hansen 2012). We also know that parents with an education are more likely to vote. If trickle-up socialization drives the effect, we would expect the effect to be strongest among parents whose children are currently in high school. If the effect is driven by parents who want to pass on

Table S4: Turnout among parents with barely ineligible children in Figure S11. N of parents with barely ineligible children in parenthesis

	Abstained last	Voted last
2013 conditioned on 2009	45.49	88.94
	(10,537)	(21,600)
2014 conditioned on 2013	19.79	69.61
	(7,419)	(25,420)

a norm of voting to their children, we might expect to see the strongest effect among parents with a high school education or more, since they are most likely to have internalized the norm of voting.

From the data I know the highest achieved education of the parents and what kind of education, if any, their children were enrolled in at the end of the latest September prior to the election. In Figure S12, I simultaneously plot the effect conditional on parents level of education and if the children were in high school. Both are measured at the end of the September before the election. As above I plot the effect in both percentage points and probits. Unfortunately, the 2015 election cannot be included in the plot as data on educational level and ongoing education was not yet updated in the public records when I acquired access[S5]. Table S5 contains turnout and group sizes.

The results show that if the child is in high school, there is no effect on the parents who themselves have a high school diploma, but a fairly large effect on parents without a high school diploma. This results applies both when I estimate the effect in percentage points and in probits. The interpretation that I suggest in the paper is that the children in high school are most likely to receive external stimuli that encourage them to vote. That means that the parents who are likely to hold a norm of voting, in other words parents with a high school diploma, does not need to step up and convey the norm of voting. At the same time parents who are less likely to hold the norm of voting, that is parents without a diploma, could become mobilized by their engaged children. In other words, they are influenced through trickle-up socialization.

For parents whose children are not in high school, we see a different pattern. For them the effects are almost identical, although the point estimate for parents who do not have a high school diploma is larger.[S6]

[S5]This part of the records is only updated up until September 2013 for the data that I have access to. Classifying children turning 18 in June 2015 based on their enrollment in September 2013 is too inaccurate to be useful as a large share of them were still in public school at that time.

[S6]It is also only the point estimate for parents without a high school diploma that excludes zero. However, as there is a large overlap between the confidence intervals for the two types of parents, I do not put too much into only one of them being different from zero.

Figure S12: Assessing heterogeneous effects for parents with and without highschool

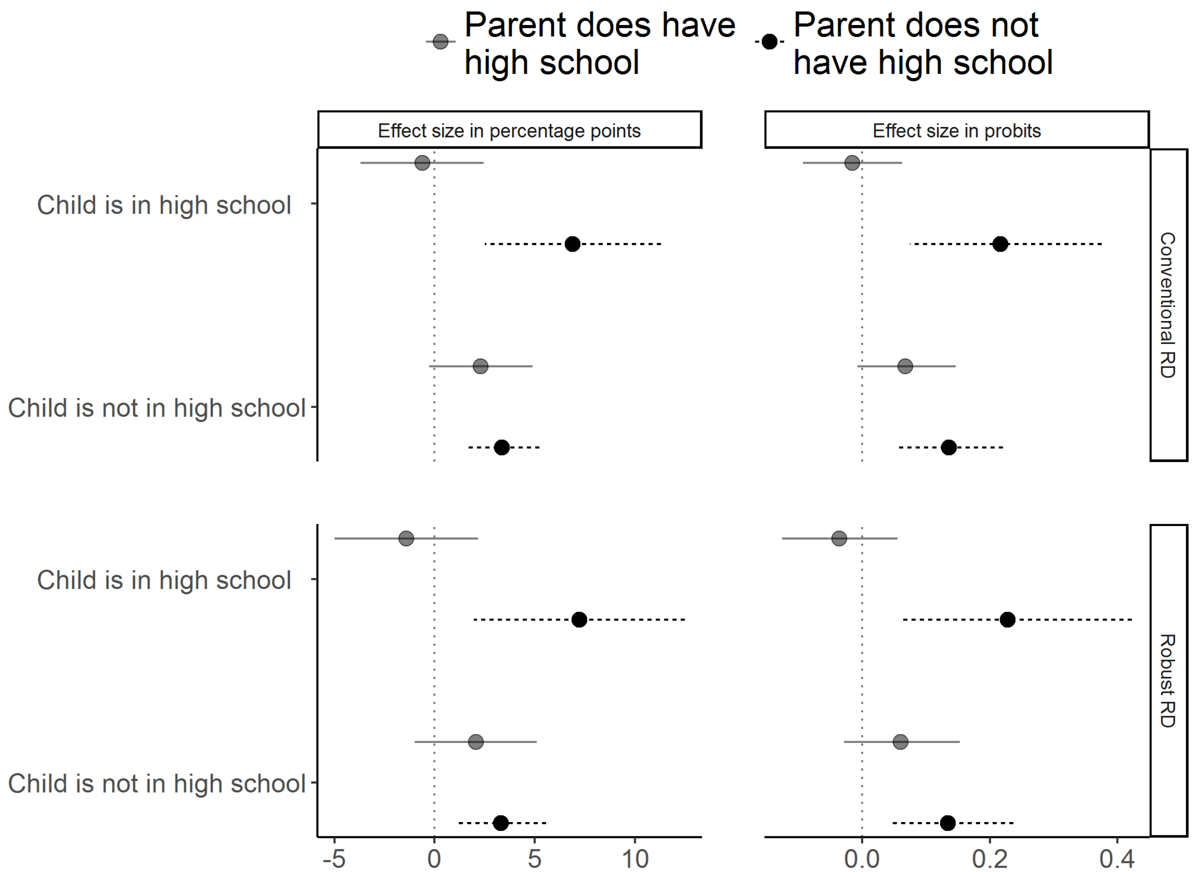

The model used to estimate the heterogeneous effects differs from the model in the main analysis in that it uses data aggregated by three days. The aggregation is necessary to assure individual anonymity in the replication data. The results are robust to using daily estimates (not reported).

For the parents with a high school diploma, we might see a stronger effect when the child is not in high school compared to when it is, because the parents feel a need to step up and convey a norm of voting. For the parents without a high school diploma, we may see a weaker effect because they are less likely to be exposed to trickle-up socialization as their children are less likely to be mobilized by other factors such as friends, teachers, and curricula.

To be clear, even though I certainly see the effects among parents in Figure S12 as indicative of both mechanisms, the differences in effects are *not* causally identified. For instance, we cannot rule out that we see the heterogeneity among parents with children taking a high school education because the parent without a high school education who want to convey a norm of voting also want their children to get a high school education. However, in conjunction with the evidence above, there are several pieces of evidence that point to both the trickle-up mechanism and parents conveying a norm of voting.

Table S5: Turnout among parents with barely ineligible children in Figure S12

	Turnout	N
Child in high school, parent with high school	81.88	(39,189)
Child in high school, parent without high school	69.91	(42,000)
Child not in high school, parent with high school	71.21	(10,164)
Child not in high school, parent without high school	55.81	(37,095)

S5 Additional heterogeneous effects

So far I have estimated heterogeneous effect in the paper, and in the Supporting Information, to tease out the mechanism. In this last section of the Supporting Information, I estimate some additional heterogeneous effects. The main effects reported in the paper are strong and robust, which is an ideal vantage point for an exploratory study of heterogeneous effects. In Figure S13, I plot both the conventional and robust effect estimates conditioned on age, income, educational level, ethnicity, if the child close to the cutoff is the oldest of the parents, and sex, interacted for parents and children. In the left panel, the effect is plotted in percentage points, in the right in probits. The probit estimates are included since the baseline turnout varies dramatically between some of the groups. Some groups, such as parents with a degree, have impressive turnout rates, and therefore a percentage point increase in their turnout rates is harder to achieve, which would be reflected in the probit estimates. It should be noted that the underlying models diverge from the models used for the main analyses, as voters are aggregated by three days or one week to ensure that no individual level information can be identified from the data. The results from the main analysis are robust to similar changes, and thus the moderations are expected to have miniscule consequences for the results.

Upfront, it is unclear what to expect. For instance, in families with immigrant backgrounds, many children might be exposed to the political system and norms as they attend school, while parents have low turnout rates. Under such conditions, children might directly influence the parents and encourage them to participate, which speaks for the effect being greater here (Wong and Tseng 2008, Bloemraad and Trost 2008). On the other hand, if parents who choose their modal behavior in a higher turnout election when their child comes of age drive the effect, we might see smaller effects. Immigrants have lower rates of participation and may not be motivated by the suggested mechanism.

A similar argument can be made for the split on income, education, and to some extent age, since turnout increases with age until voters are somewhere in their sixties (Bhatti, Hansen and Wass 2012). Older, better educated, and higher income parents are, in general, more likely to vote. Thus, one can make the argument that the effect should be weaker among these groups, since the young voters in the other homes are more likely to apply external pressure experienced outside the home to the parents. On the other hand, as suggested in the paper, high-propensity voters may drive effects because they are keen on inculcating the norm of voting in their children. Under this scenario, the effect ought to be largest among parents with a high propensity to vote. I stay agnostic with regards to the expectations, and emphasize the analyses' exploratory nature.

Based on the results in Figure S13, age does not seem to predict the effect size. The most conspicuous interaction effect is that of income, with parents earning above the median income being most likely to be mobilized. Likewise, the point estimate of the effect size is larger for parents who have a degree, although

Figure S13: Assessing heterogeneous effects

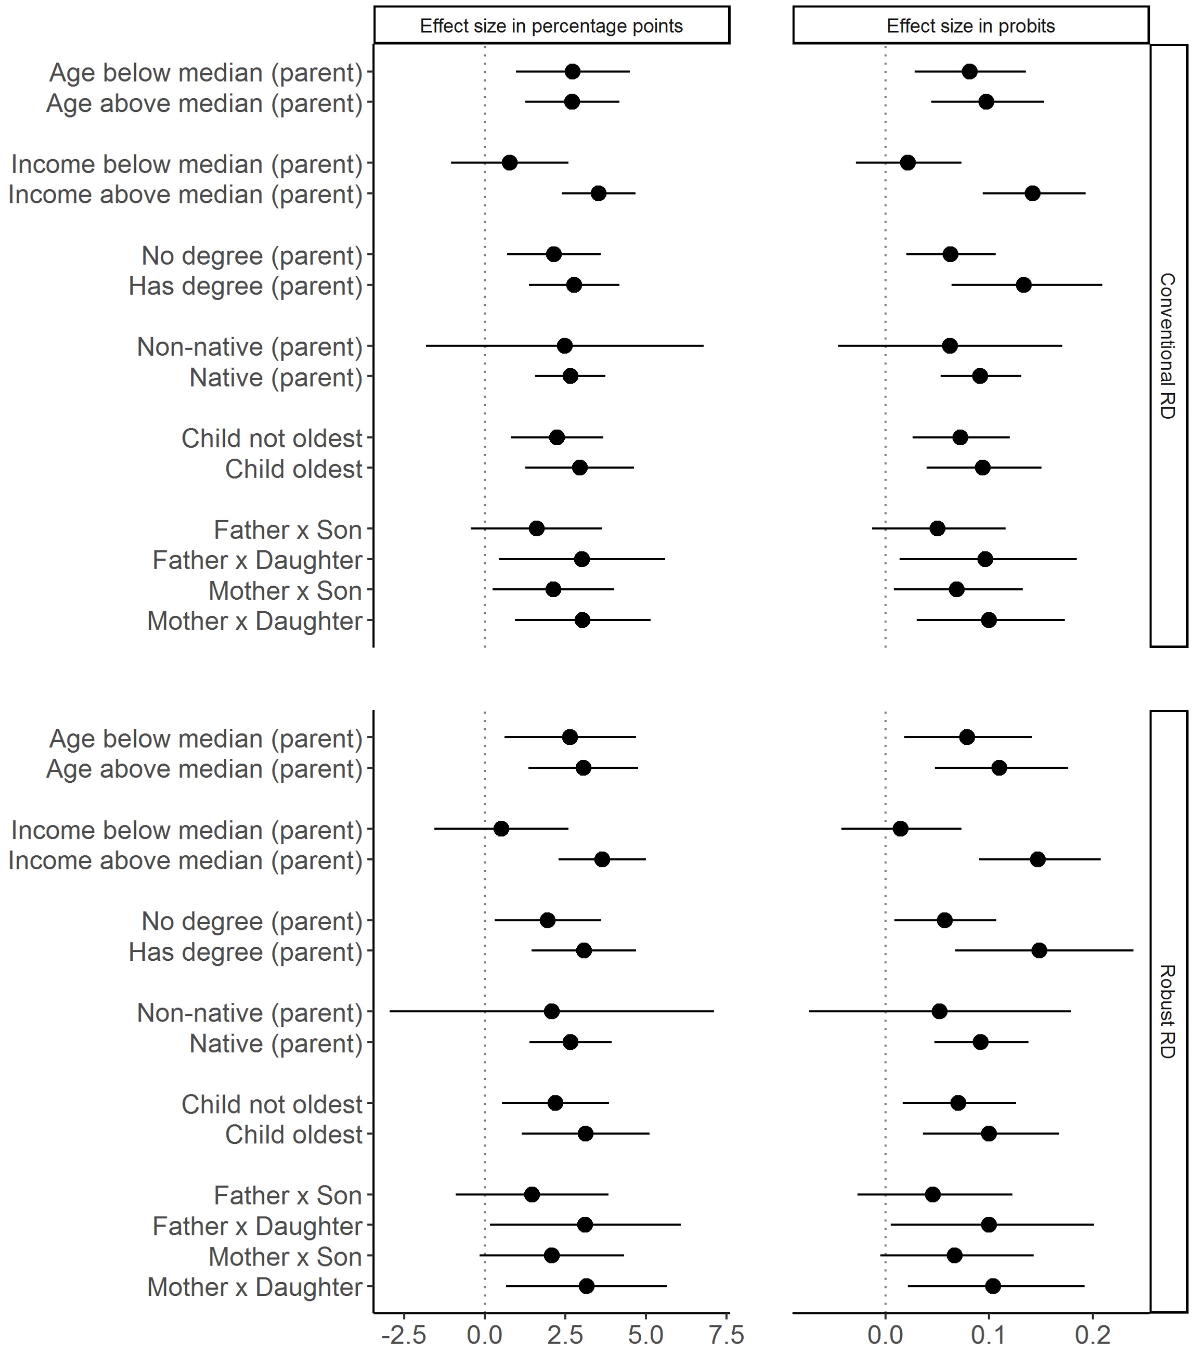

The figures contain point estimates and 95 percent confidence intervals for the precision weighted average effect for each subgroup across elections.

the difference is not as stark and there is a large overlap between the confidence intervals. When we split the effect on native and non-native parents, the confidence interval for non-natives becomes too wide for us to say anything meaningful. The point estimate is slightly larger when the child is the first-born, but we cannot rule out that this is only due to sampling variability.

Looking at the sex combinations of parents and children one things stand out: for both mothers and fathers, the effect is larger if the child coming of age is a daughter. One suggestion is that daughters are more likely to choose to vote and influnce their parents to do the same (for a theoretical argument along these lines, see Rolfe 2012).[S7] Another suggestion is that parents go to greater efforts to demonstrate good citizenship in front of their daughters, which would be consistent with experimental evidence that parents are less likely to cheat (act uncivilly) in front of daughters (Houser et al. 2015). Table S6 contains turnout and group sizes across elections for groups of parents whose children were about to come of age.

[S7]Daughters are in general more likely to turn out, than sons are.

Table S6: Turnout among parents with just ineligible children in Figure S13. N in parenthesis

	2009	2013	2014	2015	Pooled
Age below median (parent)	63.79	72.23	52.62	86.59	70.52
	(16,637)	(32,732)	(17,065)	(22,527)	(88,961)
Age above median (parent)	74.58	79.60	64.42	91.01	78.67
	(15,588)	(30,482)	(15,944)	(21,281)	(83,295)
Income below median (parent)	60.77	68.57	48.81	83.76	67.19
	(16,023)	(31,449)	(16,400)	(21,763)	(85,635)
Income above median (parent)	77.15	82.92	67.70	93.65	81.65
	(16,202)	(31,765)	(16,609)	(22,045)	(86,621)
No degree (parent)	64.50	71.79	50.69	85.65	69.87
	(23,506)	(46,144)	(22,854)	(30,101)	(122,605)
Has degree (parent)	81.16	86.56	75.48	95.51	85.81
	(8,719)	(17,070)	(10,155)	(13,707)	(49,651)
Non-native (parent)	47.07	52.20	27.91	64.19	49.61
	(4,447)	(7,908)	(2,537)	(3,153)	(18,045)
Native (parent)	72.52	79.15	60.85	90.64	77.37
	(27,778)	(55,306)	(30,472)	(40,655)	(154,211)
Child not oldest	69.70	75.69	58.44	88.79	74.73
	(17,118)	(34,884)	(17,746)	(24,320)	(94,068)
Child oldest	68.22	75.89	58.17	88.68	74.14
	(15,107)	(28,330)	(15,263)	(19,488)	(78,188)
Father x Son	69.80	75.56	59.68	88.1	74.69
	(15,992)	(33,289)	(16,455)	(22,063)	(87,799)
Father x Daughter	70.26	75.92	60.48	88.14	75.08
	(15,459)	(30,949)	(15,427)	(20,921)	(82,756)
Mother x Son	70.46	78.31	60.17	90.23	76.49
	(16,801)	(34,725)	(17,044)	(23,033)	(91,603)
Mother x Daughter	71.56	79.08	61.93	90.25	77.31
	(16,097)	(32,366)	(16,035)	(21,764)	(86,262)

The rules of eligibility for non-natives varies between the elections. For that reason, the groups are not comparable.

S6 Municipalities with data in each election and all elections

In this section, I map the municipalities that contributed with data in each year and visualize which municipalities are included in each year. I also look at which municipalities are included in all years and, accordingly, when I estimate the effect among municipalities with data in all years in table 2 of the paper.

First, I show in Figure S14, for each election year, maps of Denmark where the municipalies are colored in black if they are in the data and gray if they are not. In all years we see substantial geographic variation. In the left panel of figure S15, I show which municipalities I have data for in all four elections (black). Again, these are scattered all over the country. In the right panel of Figure S15, the municipalities are colored by the number of years in which I have data for them.

Finally, I show for each year which municipalities I have data for in Figure S16. In the figure, I draw a line for each municipality. The lines are interrupted in the years in which I do not have data. I draw black lines for the municipalities with data in all years and gray lines for the other municipalities.

Figure S14: Municipalities with data each year

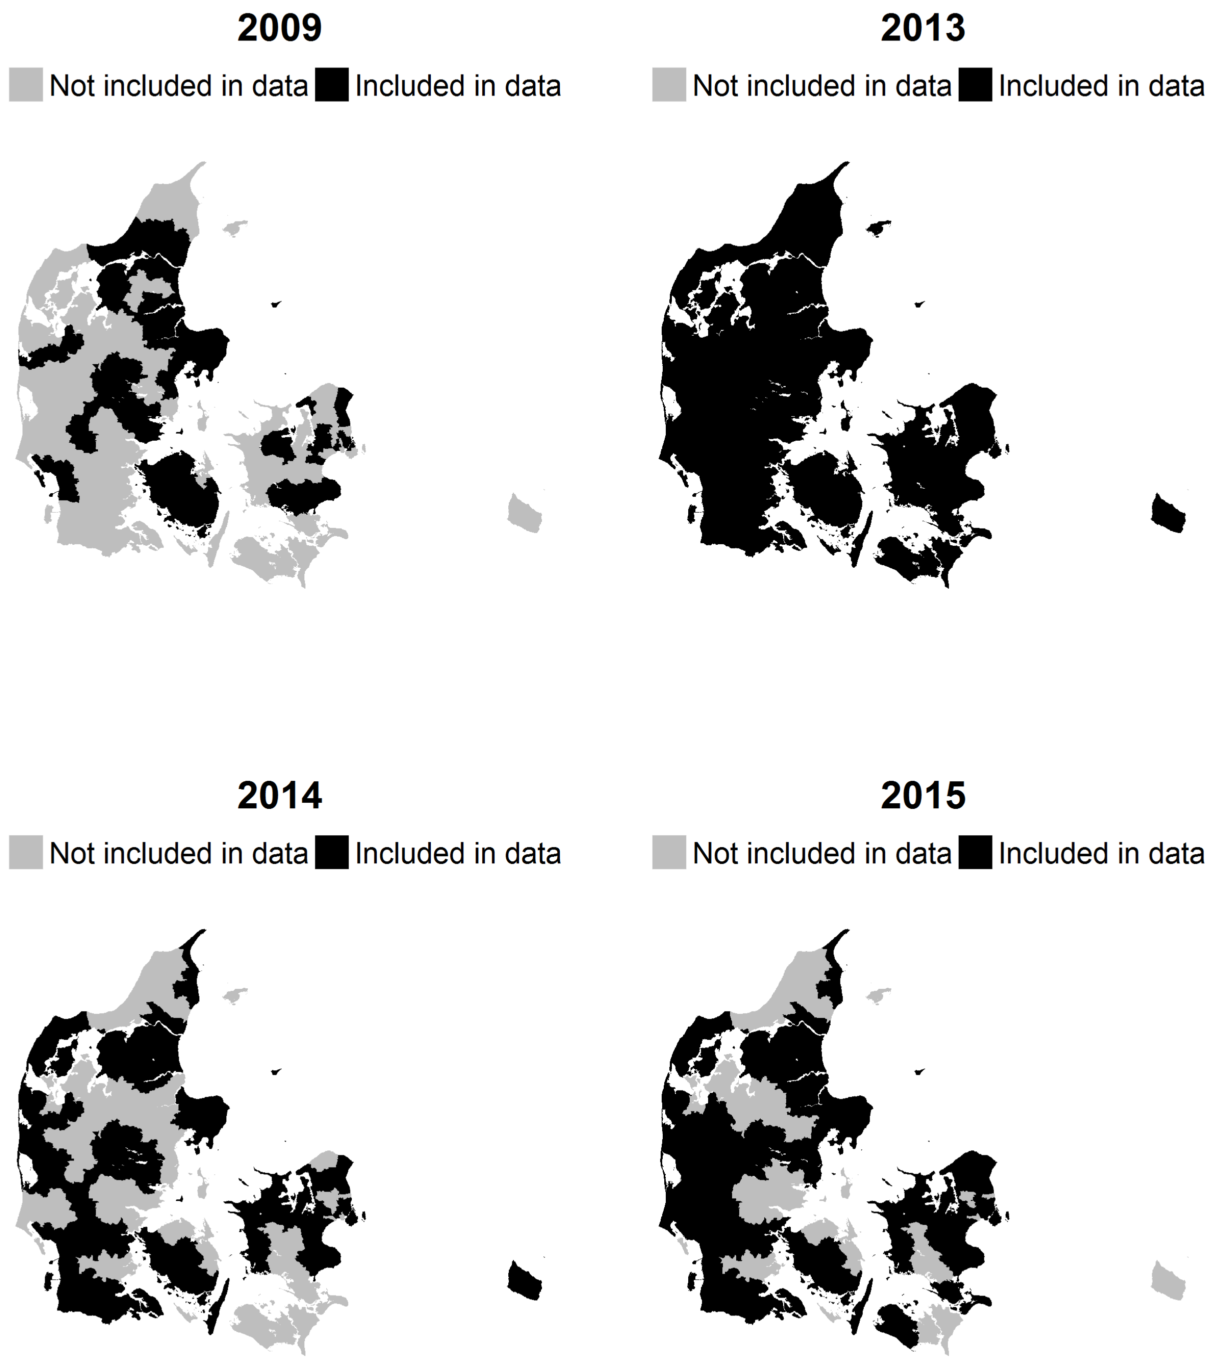

Figure S15: Municipalities with data each year

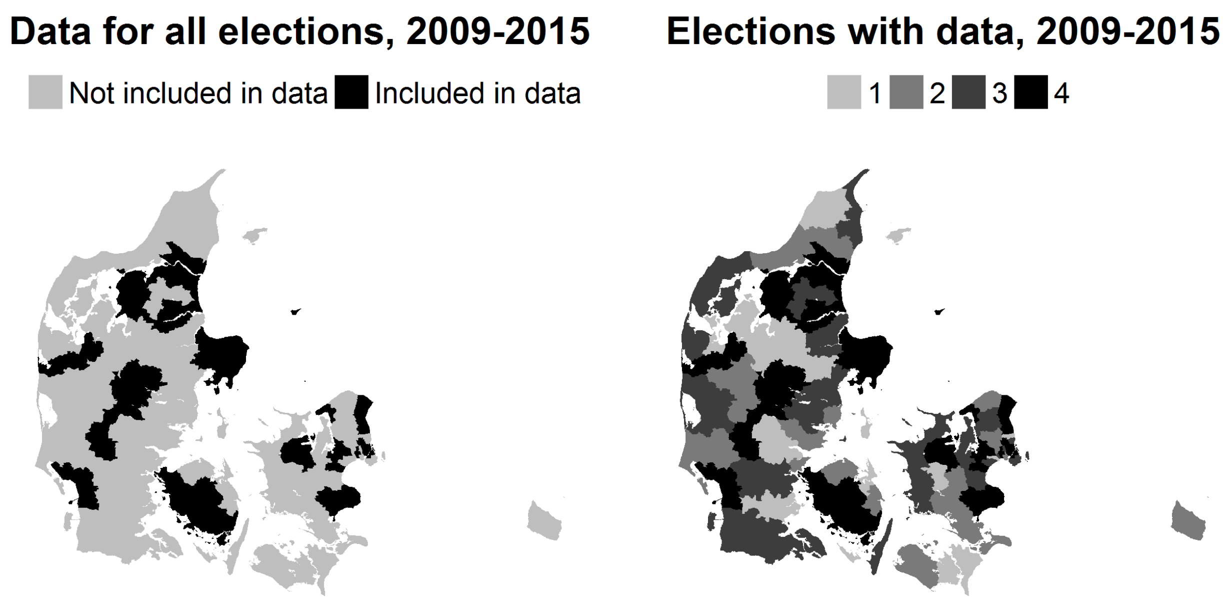

Figure S16: Overview of included municipalities

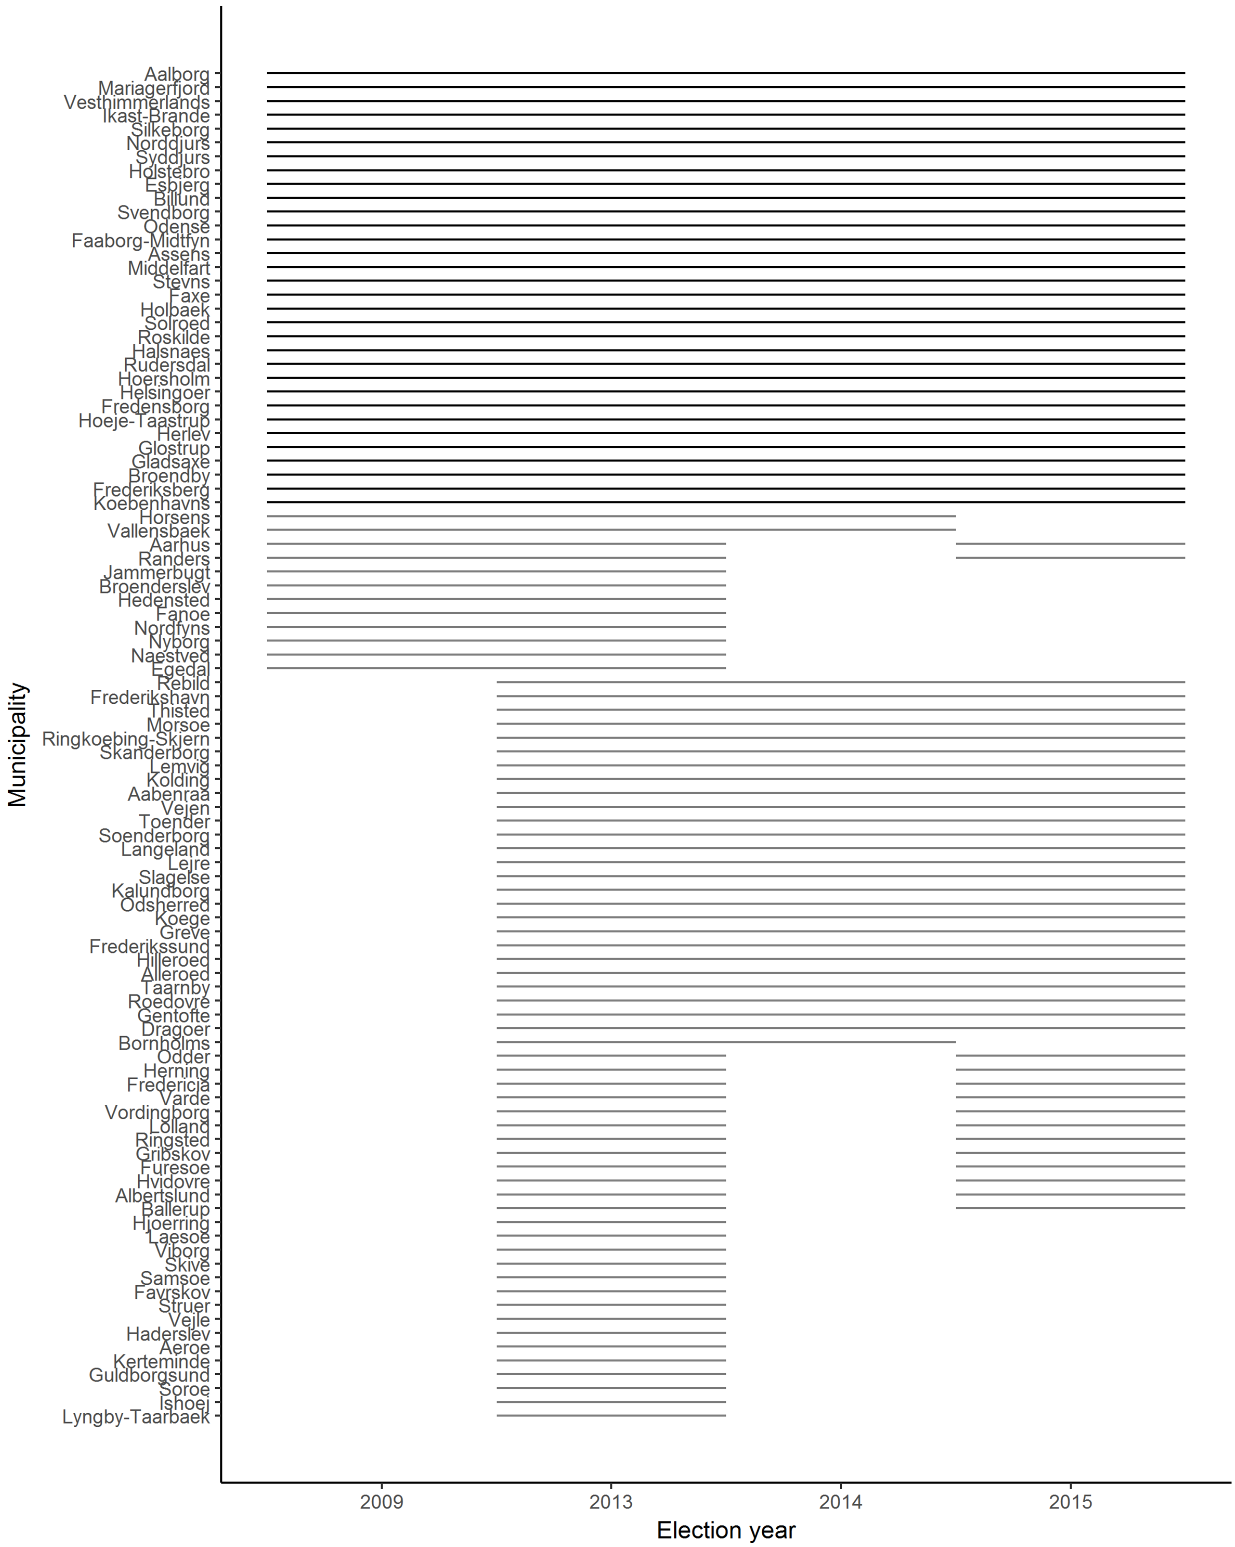

Each line segment represents an election year in which there is data for the given municipality

References

Bhatti, Yosef and Kasper M Hansen. 2012. "Leaving the nest and the social act of voting: turnout among first-time voters." *Journal of Elections, Public Opinion & Parties* 22(4):380–406.

Bhatti, Yosef, Kasper M Hansen and Hanna Wass. 2012. "The relationship between age and turnout: a roller-coaster ride." *Electoral Studies* 31(3):588–593.

Bloemraad, Irene and Christine Trost. 2008. "It's a Family Affair Intergenerational Mobilization in the Spring 2006 Protests." *American Behavioral Scientist* 52(4):507–532.

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell and Rocıo Titiunik. 2017. "rdrobust: Software for Regression Discontinuity Designs." *The Stata Journal* 17(2):372–404.

Gelman, Andrew and Adam Zelizer. 2015. "Evidence on the deleterious impact of sustained use of polynomial regression on causal inference." *Research & Politics* 2(1):2053168015569830.

Gelman, Andrew and Guido Imbens. 2017. "Why high-order polynomials should not be used in regression discontinuity designs." *Journal of Business & Economic Statistics* (just-accepted).

Houser, Daniel, John A List, Marco Piovesan, Anya Savikhin Samek and Joachim Winter. 2015. On the Origins of Dishonesty: From Parents to Children. Technical report National Bureau of Economic Research.

Linimon, Amy and Mark R Joslyn. 2002. "Trickle up political socialization: The impact of kids voting USA on voter turnout in Kansas." *State Politics & Policy Quarterly* 2(1):24–36.

McDevitt, Michael. 2006. "The partisan child: Developmental provocation as a model of political socialization." *International Journal of Public Opinion Research* 18(1):67–88.

McDevitt, Michael and Spiro Kiousis. 2006. "Experiments in political socialization: Kids Voting USA as a model for civic education reform." *The Center for Information and Research on Civic Learning and Engagement, Working Paper* 49:12.

McDevitt, Michael and Steven Chaffee. 2000. "Closing gaps in political communication and knowledge effects of a school intervention." *Communication research* 27(3):259–292.

McDevitt, Michael and Steven Chaffee. 2002. "From top-down to trickle-up influence: Revisiting assumptions about the family in political socialization." *Political Communication* 19(3):281–301.

Rolfe, Meredith. 2012. *Voter turnout: A social theory of political participation.* Cambridge University Press.

Simon, James and Bruce D Merrill. 1998. “Political socialization in the classroom revisited: The Kids Voting program.” *The Social Science Journal* 35(1):29–42.

Wong, Janelle and Vivian Tseng. 2008. “Political socialisation in immigrant families: Challenging top-down parental socialisation models.” *Journal of Ethnic and Migration Studies* 34(1):151–168.