Online Appendix to "Granting Immigrants the Right to Vote in National Elections: Empirical Evidence from Swedish Administrative Data"

Contents

A Data Availability Statement

In this paper, we apply data from Swedish registers. There are several rules and regulations on how to process and store such data. Therefore, we run our empirical analysis through a secured remote desktop system where the data are stored at a server. As a consequence, we cannot make the data available online.

Should a reader wish to gain access to these data in order to replicate our analysis, there are two ways to do so. All of the data we have used are available at Statistics Sweden (SCB) and can be ordered from them (please follow this link: https://www.scb.se/vara-tjanster/bestalla-mikrodata/). Before such a process of ordering data can begin, however, one must seek approval from the Ethical Review Board. Other researchers will then be able to order these data directly from SCB.

Another possibility would be for a person to become part of our research team temporarily. He or she would then be able to replicate our analysis by using the same remote desktop system with which we have worked. This feasibility of this option depends on where the researcher in question is based, since there are geographical restrictions on data access. If the reader is interested in this option, he or she should contact us beforehand so that we may add him or her temporarily to our research group, whereupon we shall inform the Swedish Ethical Review Board to this effect.

B Voting in Swedish elections

Individuals do not have to register to vote in Sweden; rather, all eligible individuals receive a voting rights certificate by mail three to four weeks before election day. This certificate contains basic information about the election, such as time and place, and specifies the elections in which the individual in question is eligible to vote.

At the voting station, each voter receives a number of envelopes corresponding to the number of elections they are eligible to vote in. A citizen thus receives three envelops for the national election and the two local elections (municipal and county elections), while a non-citizen eligible to vote in the local elections only receives two envelopes. Color-coded ballots for each election are sealed in the envelopes before they are handed into the election officials. Voters are under no circumstances obligated to participate in all elections they are eligible for. In fact, close to 2 percent of citizens selectively abstained in at least one election in 2010, and citizens with an immigrant background are less likely to selectively abstain (Dehdari et al. 2021).

C Details on the empirical framework

There are two traditional approaches in the RD literature on how to implement the design. The first approach consists of using all observations and include polynomials of various degrees in order to estimate the discontinuity at the cut-off. The second approach is non-parametric, which focuses on observations close to the cut-off where there is local randomization. In this study, we adopt the second approach and specify a local linear regression on each side of the cut-off when zeroing in on the observations close to the eligibility threshold (Hahn et al. 2001). Including polynomials of higher order has been criticized in Gelman and Imbens (2019), which is one reason why we instead use the non-parametric local approach.

The non-parametric approach to RD estimation involves a trade-off between bias and precision. Theoretically, one would like to zoom in on observations just around the cut-off where there is local randomization. This minimization of bias, however, comes with poor precision as one would only use a very small number of observations just around the cut-off. As a result, data-driven choices of bandwidths for the running variable in order to balance bias and precision have been presented as an alternative (Calonico et al. 2014a,b, 2017). The obstacle in our case, however, is that the institutional setting summarized in Figure 2 puts constraints on the implementation of data-driven (optimal) bandwidths. Calculating a symmetrical mean square errors (MSE) optimal bandwidth for the entire sample yields a bandwidth of 37 days, whereas the MSE-optimal bandwidths are 42 days for 1994 and 45 days for 2010 separately. For the 1994 election, the bandwidth to the right of the cut-off cannot be wider than 79 days, as we would otherwise capture the salience effect and the pure citizenship effect simultaneously. For the same reason, the bandwidth cannot be wider than 30 days to the right of the cut-off when it comes to the 2010 sample. The MSE-optimal bandwidth for 2010 is thus simply too wide. In order to be as transparent as possible, while still acknowledging the institutional constraints, we display the estimated RD coefficients for all possible bandwidths for each year and treatment window separately in the result section.

For the main analysis, we estimate local linear regressions with a uniform kernel. We believe that the first choice is the standard in the RD literature, whereas the latter choice may need some further justification. The logic behind using a triangular kernel is closely related to the core idea of the RD design, which relates to the as-if random nature of whether a unit gets assigned to either side of a specified cut-off. However, due to the administrative processes at the Migration Agency and the Tax Agency, there is increased uncertainty regarding whether a new citizen actually gains citizenship in time to be eligible as we approach the cut-off. Due to this institutional context, the relationship between gaining citizenship before the eligibility date and actually being able to vote in the national election is less reliable close to the cut-off. Hence, a triangular kernel puts more weight on observations that contain less reliable information. In Section F, we present estimates using a triangular kernel as a robustness analysis. For this reason, we apply a uniform kernel while also excluding all individuals who gained their citizenship less than seven days prior to the relevant eligibility cut-off in our donut RD estimation. The lower limit is set to 11 days. The Rdrobust command requires more than 10 mass points when running an RD with a discrete running variable.

In our case, the running variable is indeed expressed in days and is thus discrete. Lee and Card (2008) accordingly suggest that standard errors be clustered on the running variable. However, Kolesár and Rothe (2018) argue that clustering standard errors in this way produces standard errors that actually have worse coverage properties. They suggest that we instead use smaller bandwidths and stick to heteroskedasticity-robust standard errors when we have enough observations close to the cut-off, or that we use the *honest confidence intervals* discussed in their paper. Hence, we choose to estimate heteroskedasticity-robust standard errors in the main analysis, while presenting results with alternative confidence intervals in figures E5 and E6.

The regression discontinuity design is based on the premise that the underlying conditional expectation function is continuous around the cut-off. We should thus observe balance in other observable characteristics, given that the treatment is *as if* random. Moreover, there should not be any strategic bunching around the cut-off, as such bunching could indicate self-selection. In Section D, we examine whether there is any evidence of strategic bunching on either side of the cut-off.

We also assess balance in observables, thereby further addressing the issue of selection into treatment. We do this by considering age, gender, parenthood, welfare, income in $t-1$, years of education, employment status in $t-1$, and the probability of living in a large city. We assess this balance both visually and with the aid of regressions. These results can be found in Section E. The overall conclusion is that there are no clear systematic patterns of any discontinuities for these observables. This adds credibility to our claim that there is no strategic selection of certain individuals into the treatment group.

Lastly, in Figure B1, we plot voting eligibility in the national election for different values of the running variable. As we can see, however, there is a discernible drop in voting eligibility already prior to the cut-off date in both years. The only reasonable explanation for this drop is that it is due to administrative delay. According to administrators at the Migration Agency, it usually takes several days for them to inform the Tax Agency (by standard mail) of a citizenship decision, and it takes additional time for the Tax Agency to receive the citizenship letter and then record the date of citizenship in the population registers.

Thus, even in cases where the individual is granted Swedish citizenship before the cut-off date, there is no guarantee that this information will be recorded in the population registers in time for the individual to become eligible to vote in the national election. To address this problem, we drop (see the lower graphs in Figure B1) all individuals having gained citizenship less than a week before the cut-off date from our analysis. In the regression discontinuity literature, discarding observations close to the cut-off is referred

Figure B1: Eligibility by date of citizenship

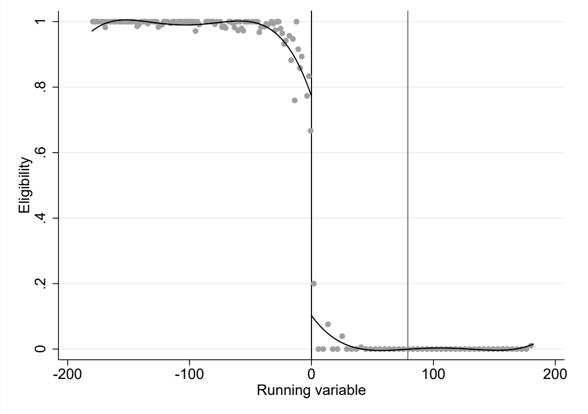

(a) Eligibility, full sample, 1994

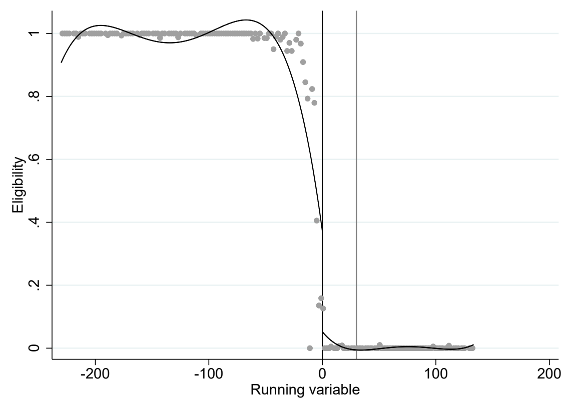

(b) Eligibility, full sample, 2010

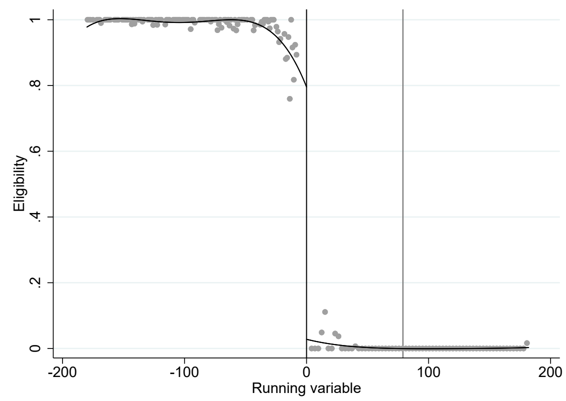

(c) Eligibility, donut sample, 1994

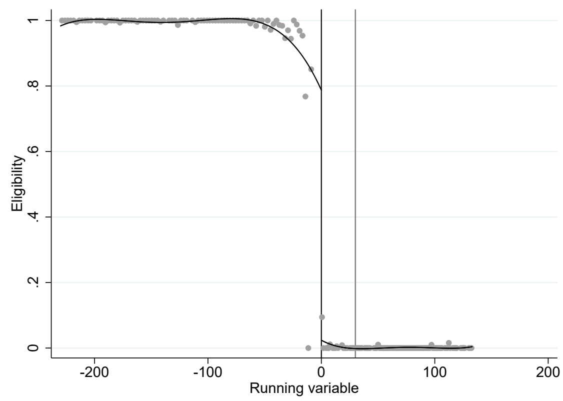

(d) Eligibility, donut sample, 2010

Notes: RD plots displaying the first-stage relationship between the running variable and being eligible to vote in national elections.

to as the *donut approach*. Our approach differs from the standard donut approach in that we only discard observations on one side of the cut-off. As noticeable, the jump in voting eligibility around the cut-off now becomes much clearer, although there is still a slight tendency for eligibility to drop already in the weeks prior to the cut-off. This is presumably because the administrative process in many cases takes more than seven days.

Even in cases where this process goes perfectly smoothly, a citizenship decision must likely be taken at least a week (5 working days) before the cut-off date for the information to be recorded in the population registers in time for the individual to be entitled to vote in the upcoming national election. In some cases, this delay will be considerably longer. In the two lower graphs in Figure B1, we show the relationship between national voting eligibility and the running variable after discarding individuals having gained citizenship the week before the cut-off.

Finally, there is a handful of individuals who are eligible to vote in the national election even though they became citizens after the cut-off date. We have not been able to find any reasonable explanation for this. However, these observations are so few that they should not affect our results. But as compliance at the threshold is not deterministic, we also estimate a fuzzy regression discontinuity model, where we let days from the eligibility cut-off serve as an instrument for actual eligibility.[12]

D Assessing the RD design: Bunching

One of the strengths of the regression discontinuity design is that it is possible to assess the local randomization at the cut-off. Figure C1 displays the number of observations for different values of the running variable around the cut-off (which is set to 0) separately for the election years 1994 and 2010.

Overall, the figures do not convey a message of complete balance over the entire period visualized. In the case of 1994 (Figure C1a), particularly few cases are decided during the month *after* the eligibility cut-off. Conversely, in 2010 (Figure C1b), there are unusually few cases just *before* the cut-off. What could give rise to these particular patterns? First, it is important to remember that the cut-off is different in the 1994 sample and the 2010 sample. For 1994, the cut-off is July 1, whereas the cut-off for 2010 is August 19. Hence, these patterns do not fit with a strategic sorting story, which would suggest that an alarmingly high number of individuals were processed *just before* the cut-offs in both years. It is also noteworthy that for 1994, when there are fewer observations after the cut-off, there does not seem to be bunching exactly around the cut-off.

We consider the most likely explanation for the results to be found in the extensive and institutionalized Swedish summer vacation in July, during which the level of activity

[12]In general, the quality of the data is very good. There are a few minor errors in the registries (e.g., duplicate observations for individuals). We drop these individuals from our analyses.

Figure C1: Assessing the RD design: Bunching histograms

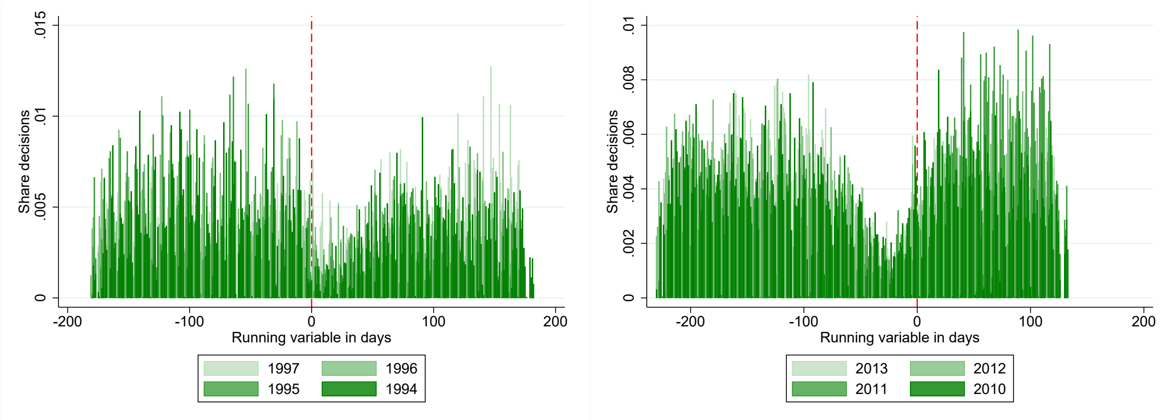

(a) 1994–1997

(b) 2010–2013

Notes: Histogram of number of individuals granted citizenship for different values of the running variable for the years 1994–1997 and 2010–2013. The solid red lines indicate cut-offs to be included in the election roll.

at any agency is low for several weeks. Note that the time periods with unusually few cases (just after the cut-off in 1994 and before the cut-off in 2010) largely overlap with the month of July, during which large segments of public administration in Sweden are far less active due to people being on vacation. To support this story further, both histograms presented in Figure C1 have been overlayed with data for the additional years for each term of office. The first histogram thus displays data from 1994–1997 and the second displays data from 2010–2013. The actual election year is displayed in bright green whereas the other years are paler. No elections occurred in these additional years, which means that no sorting due to eligibility status should occur. Yet, we observe exactly the same patterns as those seen in the election years of 1994 and 2010.

In Figure C2, we present additional overlapping histograms for the terms of office 1998–2001, 2002–2005, and 2006–2009. The cut-off has been placed at the exact date a person had to be registered as a citizen to be included in the election roll for the elections in 1998, 2002, and 2006. The main message is that we observe the exact same pattern again as we did in Figure C1. For all years in our data, we observe much fewer decided cases during the vacation month of July regardless of whether or not there is an election.

To further analyze bunching, we have also run a formal density test. There are some obstacles in terms of how to design such a density test for our setting. First, the running variable is discrete and the standard density tests are designed for continuous running variable (McCrary 2008). Second, we need to take into account that the month of July stands out due to vacations. We have designed a test where we use all data available from 1994 to 2013. This includes a total of 5 terms of office. We chose the same cut-off as the preceding election as the placebo cut-off in the subsequent years during a term of office. We then calculate the total number of cases decided during each year. Next, we calculate the share of citizenship decisions determined on a given calendar day in each year. We create a dummy variable that equals 1 for election years and otherwise 0. We then run a regression with the share of decided cases as the dependent variable and the election year dummy as the variable of interest. We include term-of-office-fixed effects and calendar-day-fixed effects. Hence, the comparison will be made in a term of office (the same cut-off) and the same calendar date. We then restrict the analysis to the same window (11–79 days) around the cut-off that we use for the main RD analysis. If sorting did occur around the cut-off in election years, the estimated coefficient will turn out to be statistically significant. We find the opposite in Figure C3. The estimated coefficients are stable, small, and statistically insignificant for the large majority of all specifications. In conclusion, the pattern in decided cases observed in election years does not deviate from the standard pattern seen in all other years in the data.

Figure C2: Assessing the RD design: Additional bunching histograms

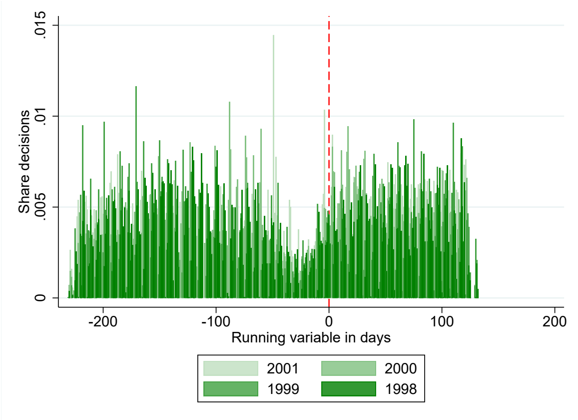

(a) 1998–2001

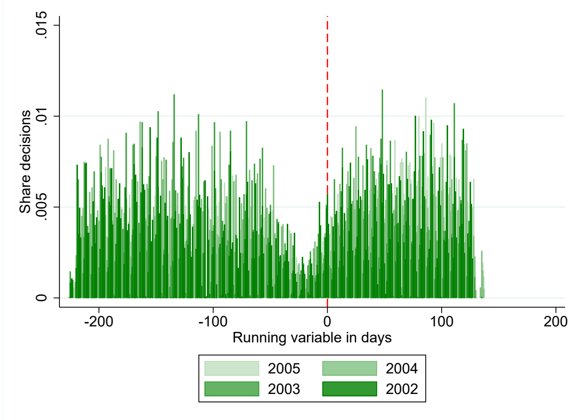

(b) 2002–2005

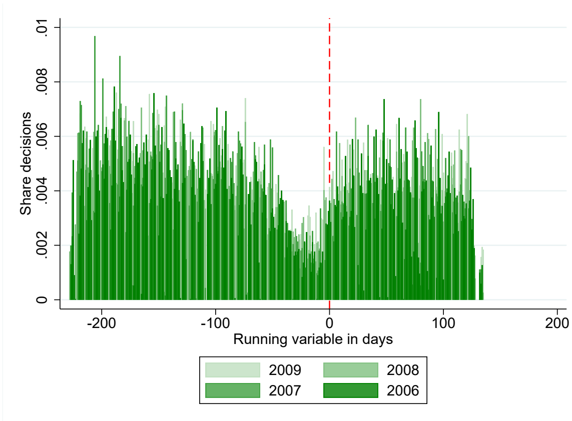

(c) 2006–2009

Notes: Histogram of number of individuals granted citizenship for different values of the running variable for the years 1998–2001, 2002–2005, and 2006–2009. The solid red lines indicate cut-offs to be included in the election roll.

Figure C3: Density test

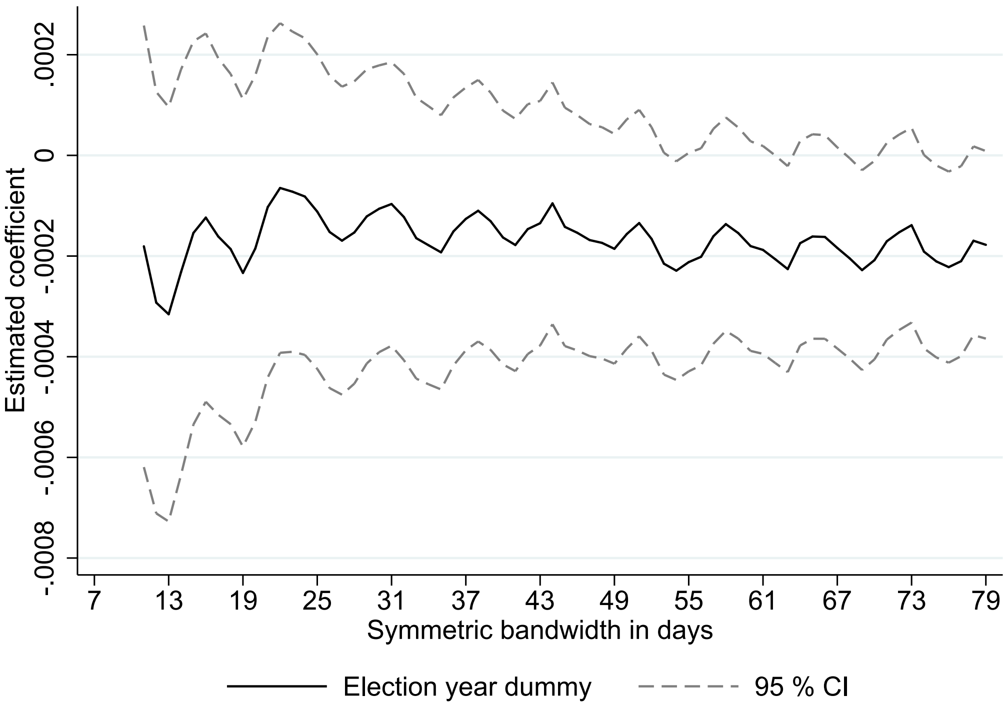

E Balance tests

Figures D1 and D2 show balance tests, including age, income ($t-1$), employment ($t-1$), parental status, years of education, belonging to a household on welfare, and living (or not) in a major city. The running variable measures days before/after the cut-off. As we can see, there is no clear jump in any of the covariates at the threshold. The small "jumps" that we do see in some figures do not translate into any substantial difference in the y-variable.

We also show the balance more formally using local linear regressions and a uniform kernel. The results, seen in Table D1, demonstrate that almost none of the covariates jump significantly at the eligibility cut-off. Two exemptions are employment and welfare, but neither of these is likely a cause for concern. First, we run 42 balance tests in Table D1. A couple of estimates will be statistically significant is to be expected as a result of pure chance. Second, the point estimates are very different when comparing 2010 to 1994. Third, in the case of employment, the effect is slightly negative in 2010. This would suggest that employed individuals were slightly *less* likely to end up with the right to vote in all elections in 2010. If anything, this would depreciate our estimates, as we expect employed individuals to vote more than non-employed individuals. In the case of welfare, we again observe no significant changes in 1994, but a negative drop in 2010. To rule out that this matters for our main estimates, we have rerun the baseline results with controls in Figure E2 and Figure E3, which does not matter for the outcome.

In addition to the individual characteristics presented so far, we also consider region of birth. Figure D3 plots the running variable against the share of new citizens who were born in Africa, Asia, Europe, or South America. We exclude citizens born in North America or Oceania due to insufficient observations. We also show local linear estimates in Table D2. Our general takeaway is that neither the plots nor the regressions give us a strong reason to suspect selection patterns. First, the plots for Africa and South America balance very well for both 1994 and 2010. For Asia and Europe, we also see small and insignificant estimates for 2010, while a couple of the coefficients do, in fact, turn out to be significant in 1994. However, there are a couple of reasons for not being alarmed by these results. First, in the case of both Asia and Europe, the estimates are inconsistent: While there is evidence of negative (positive) selection for Europe (Asia) for the two larger bandwidths, the estimates switch to positive (negative) in the smaller bandwidth. Second, in the next section, we run heterogeneity estimates suggesting that individuals from Europe and Asia respond very similarly to the treatment. Hence, even with a slight imbalance in the samples, it is unclear to what extent this would affect the results. Finally, much like in the case of socio-economic and demographic characteristics, we run a large number of regressions, where one or two are likely to come out significant by chance alone. These covariates are further included in Figure E2 and Figure E3 as already discussed above.

Figure D1: Balance plots 1

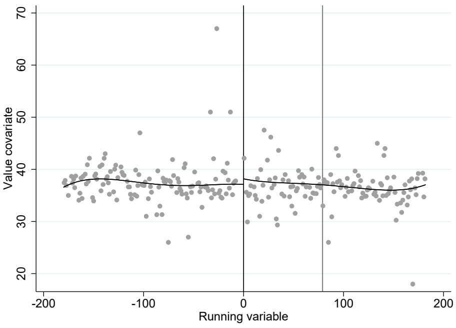

(a) Age, 1994

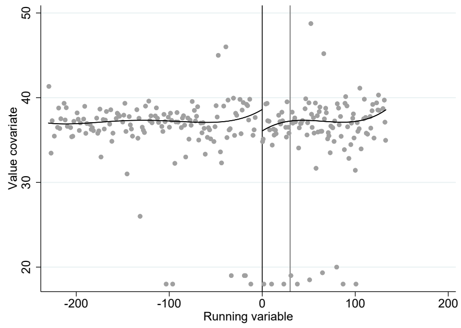

(b) Age, 2010

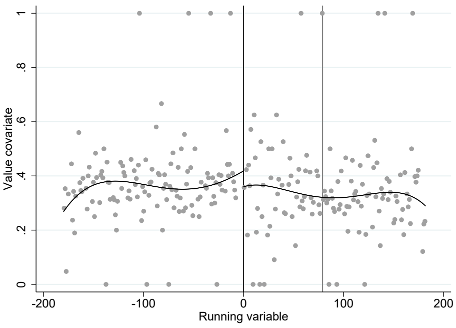

(c) Living in a large city, 1994

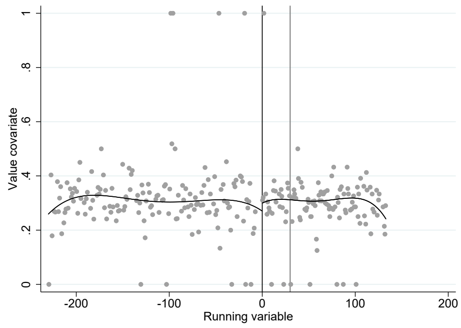

(d) Living in a large city, 2010

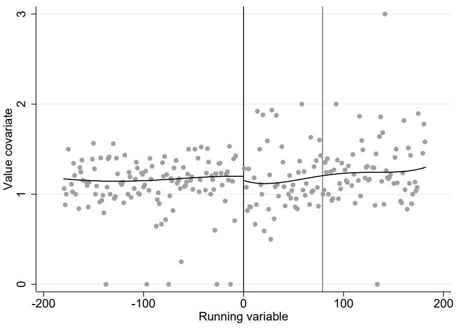

(e) Having children (binary), 1994

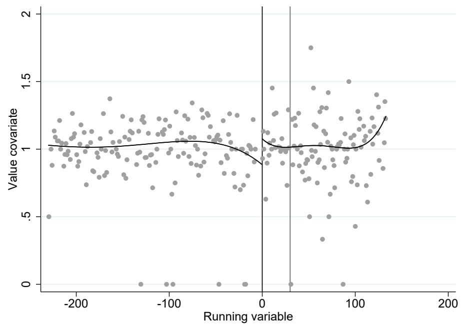

(f) Having children (binary), 2010

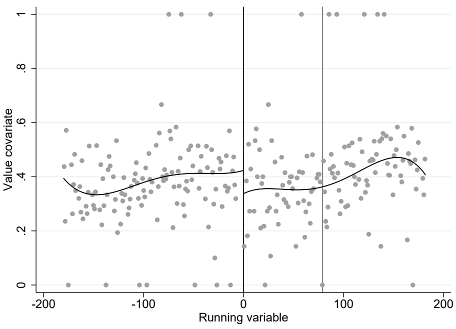

(g) Family on welfare, 1994

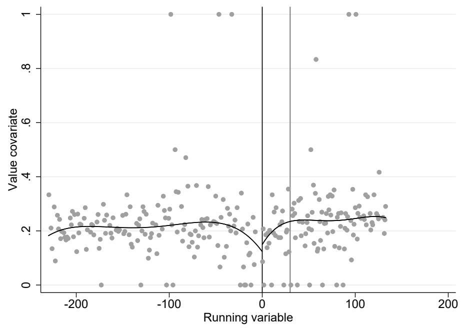

(h) Family on welfare, 2010

Notes: RD plots according to Calonico et al. (2015). The running variable, displayed on the x-axis, measures days before/after the eligibility cut-off for 1994 and 2010, respectively. The treatment window is placed in a chronological calendar order to the left in each figure. The y-axis displays the outcome.

Figure D2: Balance plots 2.

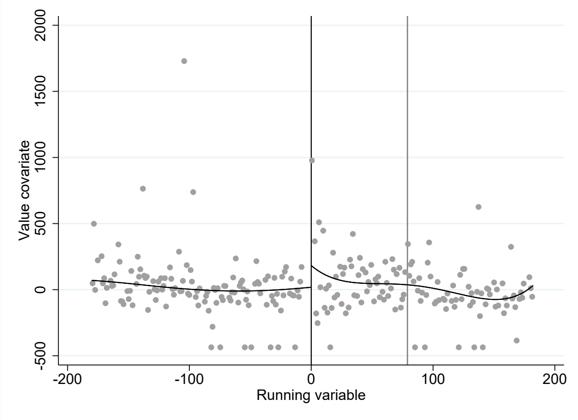

(a) Income t-1, residualized, 1994

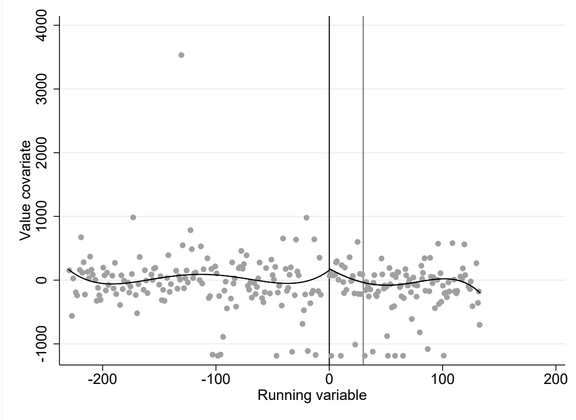

(b) Income t-1, residualized, 2010

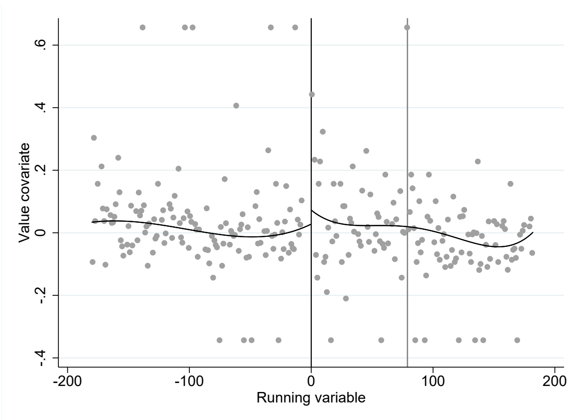

(c) Employment status t-1, residualized, 1994

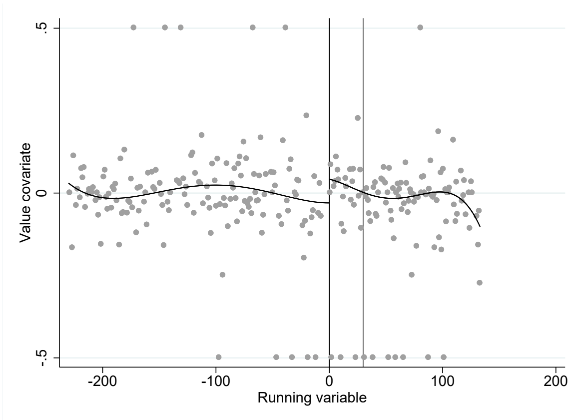

(d) Employment status t-1, residualized, 2010

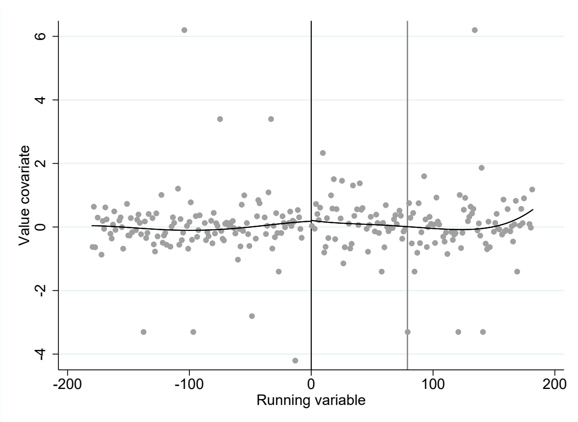

(e) Years of education, residualized, 1994

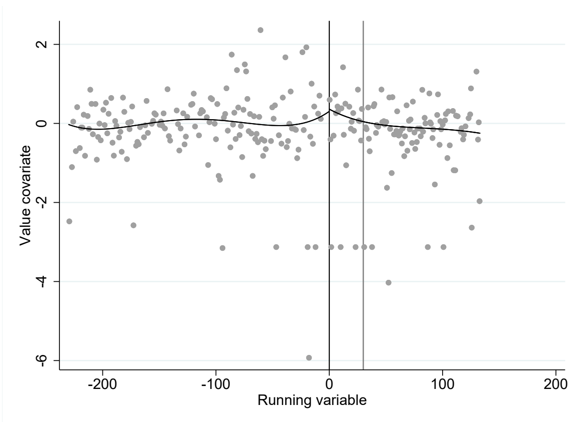

(f) Years of education, residualized, 2010

Notes: RD plots according to Calonico et al. (2015). The running variable, displayed on the x-axis, measures days before/after the eligibility cut-off for 1994 and 2010, respectively. The treatment window is placed in a chronological calendar order to the left in each figure. The y-axis displays residuals from a regression where the outcome is regressed on the year of obtaining citizenship (1994 and 2010). This is to compensate for a pure year effect (arising from differences in the inflation and the business cycle.

Table D1: Balance table

	Age	Big City	Parent	Years of Education	Income (year prior)	Employment (year prior)	Welfare
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
PANEL A: ELECTION YEAR 1994							
BW 79/79	-1.345	0.017	0.154	0.021	-28.807	-0.002	0.058
	(0.795)	(0.033)	(0.088)	(0.209)	(49.025)	(0.032)	(0.033)
Obs, below cut-off	2198	2200	2198	2198	2200	2200	2198
Obs, above cut-off	3828	3831	3830	3830	3828	3828	3830
BW 50/50	-0.352	0.055	0.057	-0.059	-4.066	0.021	0.041
	(1.081)	(0.046)	(0.121)	(0.300)	(69.067)	(0.045)	(0.046)
Obs, below cut-off	951	951	950	950	951	951	950
Obs, above cut-off	2078	2079	2079	2079	2078	2078	2079
BW 30/30	3.508	0.035	0.211	-0.123	-25.711	-0.000	0.026
	(1.577)	(0.068)	(0.172)	(0.444)	(104.645)	(0.065)	(0.066)
Obs, below cut-off	377	377	377	377	377	377	377
Obs, above cut-off	1239	1239	1239	1239	1238	1238	1239
PANEL B: ELECTION YEAR 2010							
BW 79/30	0.900	0.016	-0.124	-0.191	-205.560	-0.084	-0.006
	(0.817)	(0.033)	(0.083)	(0.259)	(121.530)	(0.036)	(0.028)
Obs, below cut-off	1837	1837	1837	1837	1833	1833	1837
Obs, above cut-off	2552	2552	2552	2552	2549	2549	2552
BW 50/30	0.516	0.005	-0.178	-0.009	-164.519	-0.088	-0.061
	(0.930)	(0.038)	(0.093)	(0.299)	(142.199)	(0.041)	(0.031)
Obs, below cut-off	1837	1837	1837	1837	1833	1833	1837
Obs, above cut-off	1198	1198	1198	1198	1195	1195	1198
BW 30/30	0.075	-0.067	-0.084	0.287	-121.169	-0.069	-0.093
	(1.117)	(0.046)	(0.108)	(0.366)	(175.313)	(0.050)	(0.036)
Obs, below cut-off	1837	1837	1837	1837	1833	1833	1837
Obs, above cut-off	635	635	635	635	634	634	635
Functional form	Linear	Linear	Linear	Linear	Linear	Linear	Linear
Kernel	Uniform	Uniform	Uniform	Uniform	Uniform	Uniform	Uniform
Covariates	No	No	No	No	No	No	No

Notes: Standard errors within parentheses. Income, employment, and years of education are residuals from a regression where the outcome is regressed on the year of obtaining citizenship in 1994 and 2010. This is to compensate for a pure year effect.

Figure D3: Balance plots: countries of origin

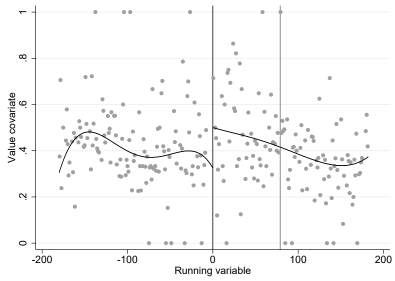

(a) Europe, 1994

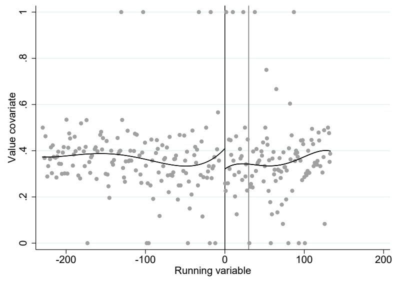

(b) Europe, 2010

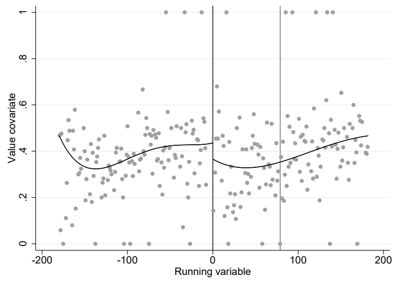

(c) Asia, 1994

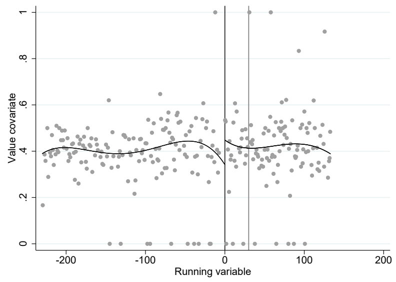

(d) Asia, 2010

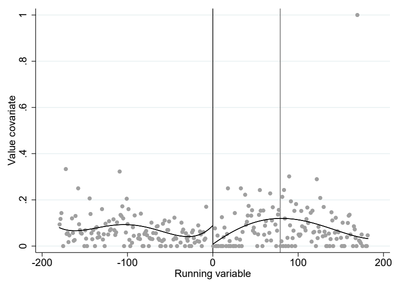

(e) South America, 1994

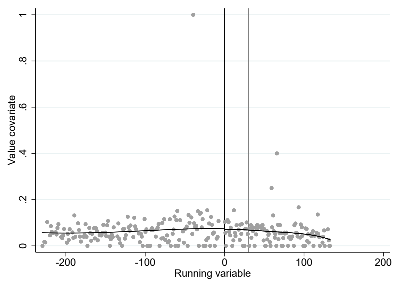

(f) South America, 2010

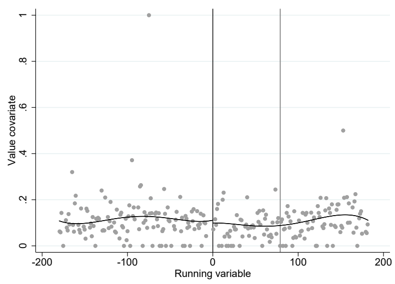

(g) Africa, 1994

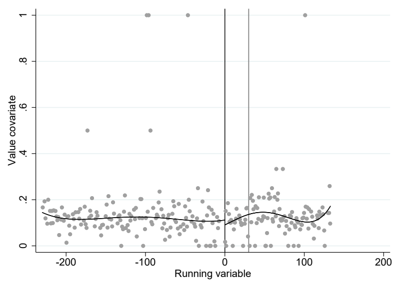

(h) Africa, 2010

Notes: See Figure D1 for more information on the RD design.

Table D2: Balance table: countries of origin

	Africa (1)	Asia (2)	Europe (3)	South America (4)
PANEL A: ELECTION YEAR 1994				
BW 79/79	0.018	0.055	-0.079	0.006
	(0.020)	(0.032)	(0.032)	(0.015)
Obs, below cut-off	2198	2198	2200	2198
Obs, above cut-off	3828	3828	3831	3828
BW 50/50	-0.005	0.119	-0.153	0.032
	(0.029)	(0.045)	(0.045)	(0.020)
Obs, below cut-off	951	951	951	951
Obs, above cut-off	2078	2078	2079	2078
BW 30/30	-0.043	-0.026	0.035	0.019
	(0.045)	(0.066)	(0.065)	(0.027)
Obs, below cut-off	377	377	377	377
Obs, above cut-off	1239	1239	1239	1239
PANEL B: ELECTION YEAR 2010				
BW 79/30	-0.019	0.017	0.001	0.012
	(0.022)	(0.035)	(0.034)	(0.019)
Obs, below cut-off	1837	1837	1837	1837
Obs, above cut-off	2550	2550	2552	2550
BW 50/30	-0.042	0.022	0.013	0.001
	(0.025)	(0.041)	(0.039)	(0.022)
Obs, below cut-off	1837	1837	1837	1837
Obs, above cut-off	1197	1197	1198	1197
BW 30/30	-0.065	0.003	0.042	0.004
	(0.027)	(0.050)	(0.047)	(0.028)
Obs, below cut-off	1837	1837	1837	1837
Obs, above cut-off	634	634	635	634
Functional form	Linear	Linear	Linear	Linear
Kernel	Uniform	Uniform	Uniform	Uniform
Covariates	No	No	No	No

Notes: Standard errors within parentheses.

F Robustness analysis for the main results

In this section, we present various robustness checks for our main RD analysis regarding the effect of having the right to vote in the more salient national election on voter turnout in local elections. To facilitate for the reader, the y-axis in all bandwidth figures in the following sections of the appendix are capped at (-1, 1) if necessary. Otherwise, the figures would be pushed together, thus making it more difficult to interpret the fluctuation of the point estimate. This choice implies that on the rare occasions that the 95% confidence interval generates margins above 1 or below -1 (i.e., larger than a 100 percentage point difference), the exact estimate of the upper and lower margins will be missing from the figures. Note that in all these cases, the point estimates are clearly not statistically significantly different from zero.

In our main analysis in Figure 4, we use a 7-day donut window when we run our regressions. In Figure E1, we check whether our results are sensitive to the specification of the donut window. Specifically, we increase the donut window from 1 to 14 days away from the cut-off. We present the results for both the sharp specification and the fuzzy one. Our conclusion from this robustness check is that our main findings are not particularly sensitive to the specification of the donut window. The estimated coefficients are fairly stable across specifications, although the confidence intervals widen for some donut specifications.

Figures E2 and E3 display the results for various alternative specifications of the main analysis presented in Figure 4. First, we run the same analysis as in the main text but include covariates. Second, we use a triangular kernel instead of a uniform one. Third, we present the results from an analysis with a local quadratic polynomial instead of a local linear specification. In line with the conclusions in Gelman and Imbens (2019), we do not estimate higher polynomials than quadratic ones. The sharp estimates are presented in Figure E2 and the fuzzy estimates in Figure E3.

The point estimates are overall stable in Figure E2 and Figure E3. When we shrink the bandwidth to a very small window, the point estimate begins to wobble and the confidence intervals become very wide, which is what we would expect given that we only include 4 days of bandwidth for the smallest specification above the cut-off (11 days bandwidth, removing 7 days in the donut). In general, the results we find are in line with our main findings. We generally estimate positive coefficients. These coefficients are often statistically significant, although not for all bandwidth specifications (especially not for the 2010 sample).

In the main text, we display the results for bandwidth sizes of 11–79 days for the 2010 election. The reason is that we want to mimic the 1994 analysis as closely as possible to facilitate comparisons. However, we cannot increase the bandwidth above the cut-off to more than 30 days for 2010, since this would mean that we also include observations in the citizenship treatment window. Hence, we fix the below cut-off bandwidth to 30 days in Figure 4 in the main text. In Figure E4, we, for the sake of completeness and transparency, display the results for symmetric bandwidth calculations for the 2010 sample below 30 days of the running variable.

In the main text, we also discuss the issue of inference and how to estimate the confidence intervals for our RD coefficients. In short, there is a discussion in the literature on how to perform inference when implementing a non-parametric RD specification with different polynomials on both sides of the cut-off. Calonico et al. (2014b), Calonico et al. (2014a), and Calonico et al. (2017) present a `rdrobust` package where they argue that one needs to separate between point estimation and inference when running a non-parametric RD estimation. To overcome this issue, they suggest estimating bias-corrected robust confidence intervals. We estimate such confidence intervals in Figure E6. The construction of these confidence intervals is based on an under-smoothed point estimate, meaning that the point estimate for a particular bandwidth is not centered in the CI. In Figure E6, we thus only plot the bias-corrected confidence intervals and not the point estimate. Kolesár and Rothe (2018) and Armstrong and Kolesár (2020) criticize these bias-corrected confidence intervals for being too wide. They instead suggest using *honest confidence intervals*. We estimate such confidence intervals in Figure E5. One obstacle when estimating *honest CIs* is that you need to choose, a priori, a bound of the second derivative of the conditional expectation function at the cut-off. The estimated confidence intervals in Figure E5 assume a linear approximation with a bound of 0 of the second derivative.

Lastly, we present the results from a robustness check in Figure E7, where we exclude individuals born in another EU member state. The rationale for this analysis is the decision taken by the European Council in December 1994 granting all EU nationals the right to vote in local elections if they resided in Sweden in the same manner as Swedes (European Council 1994). What this means is that in the 2010 election, EU nationals could have resided in Sweden for a much shorter period of time and still be eligible to vote in the local elections (although not in the national election) in comparison to those originating outside the EU, who need to be residents in Sweden for the last three years to vote in the local elections. Our prior is that these EU nationals should be evenly spread out over the values of the running variable, but it is still important to analyze whether our main results change if we exclude EU nationals. The conclusion from Figure E7 is that the findings from this sub-analysis are in line with the main findings in Figure 4 in the main text.

Figure E1: Robustness analysis specifying the donut window

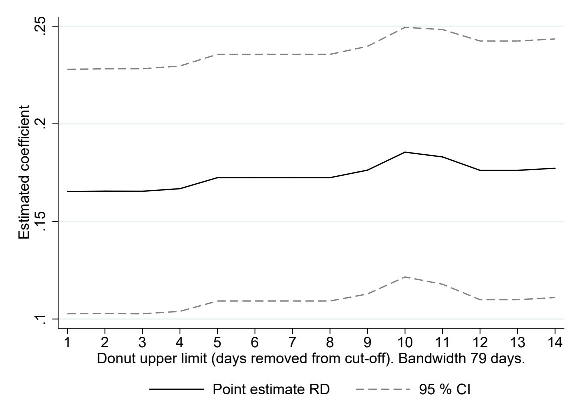

(a) Sharp estimates, 1994

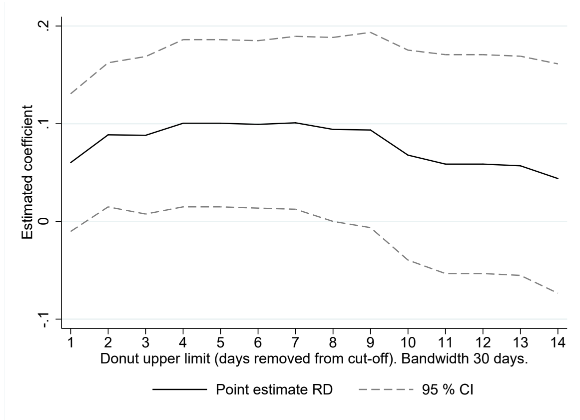

(b) Sharp estimates, 2010

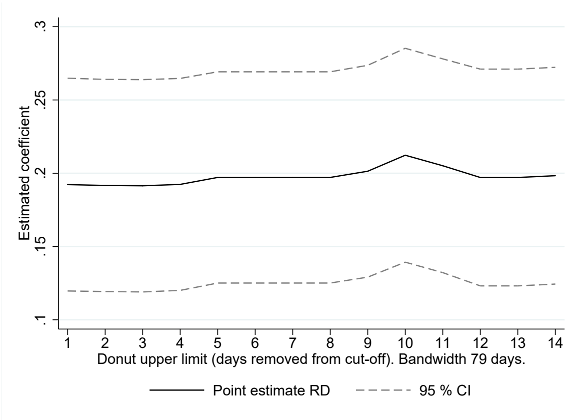

(c) Fuzzy estimates, 1994

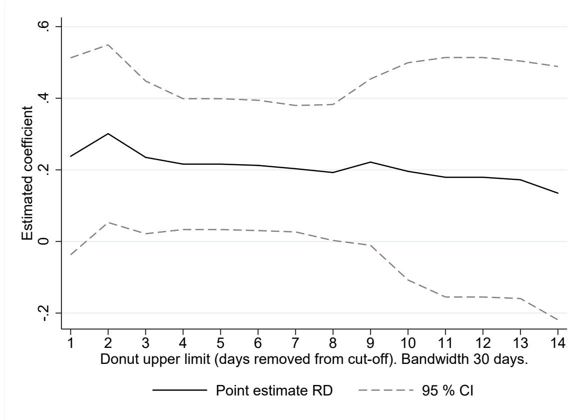

(d) Fuzzy estimates, 2010

Notes: The x-axis refers to the upper limit of the donut, whereas the lower limit is always set to 0. The RD estimation is the same as in the main analysis in Figure 4, where we include no covariates and use a uniform kernel. The bandwidth is 79 days for 1994 and 30 days for 2010.

Figure E2: Robustness analysis for covariates, kernels, and polynomials. Sharp estimates.

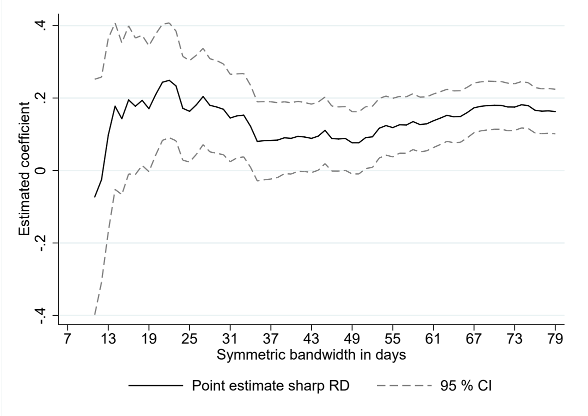

(a) With covariates, 1994

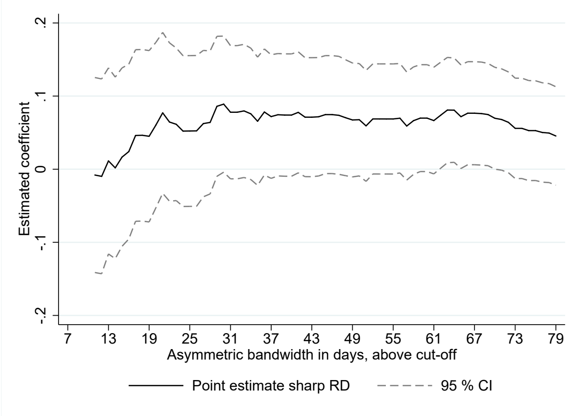

(b) With covariates, 2010

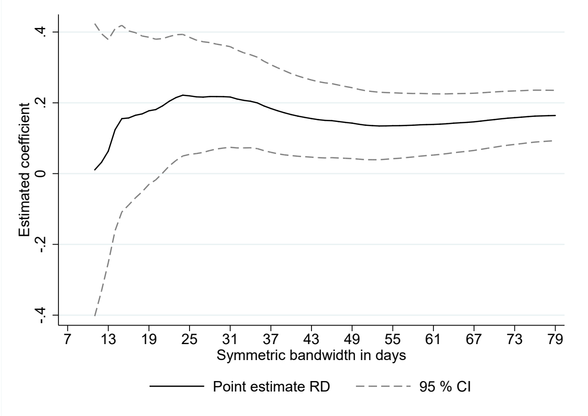

(c) Triangular kernel, 1994

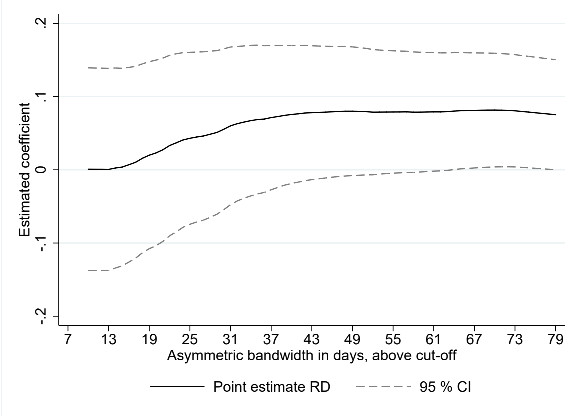

(d) Triangular kernel, 2010

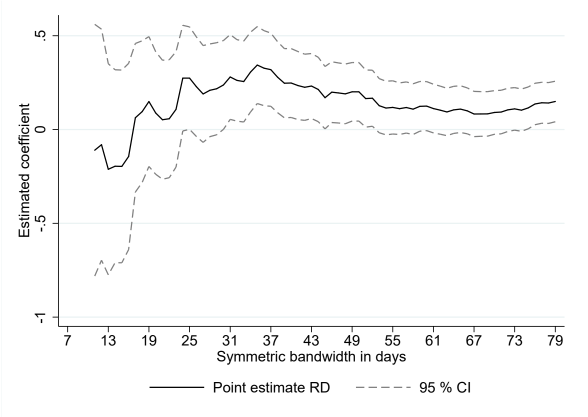

(e) Quadratic polynomial, 1994

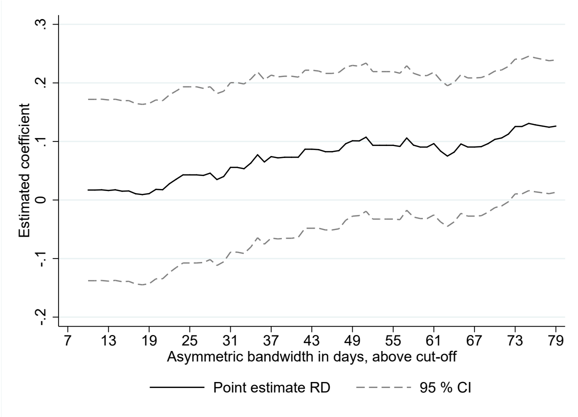

(f) Quadratic polynomial, 2010

Notes: Subfigures (a) and (b) display the main sharp RD estimates but with included covariates. Subfigures (c) and (d) employ a triangular kernel instead of a uniform one. Subfigures (e) and (f) employ a local quadratic polynomial instead of a local linear specification. For 1994, the bandwidth is symmetric, while it is asymmetric for 2010, with the upper limit always set at 30. CIs are capped at (-1,1)

Figure E3: Robustness analysis for choice of bandwidths, kernels, and polynomials. Fuzzy estimates.

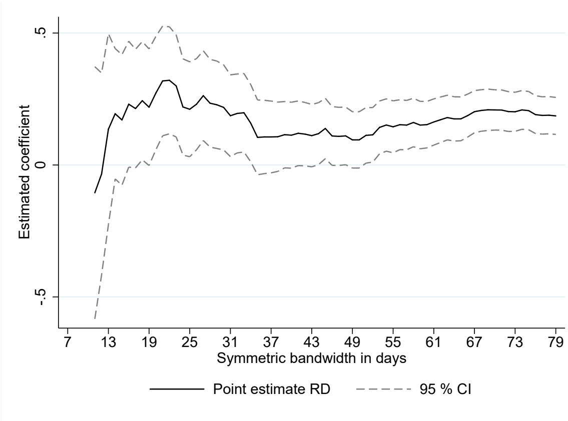

(a) With covariates, 1994

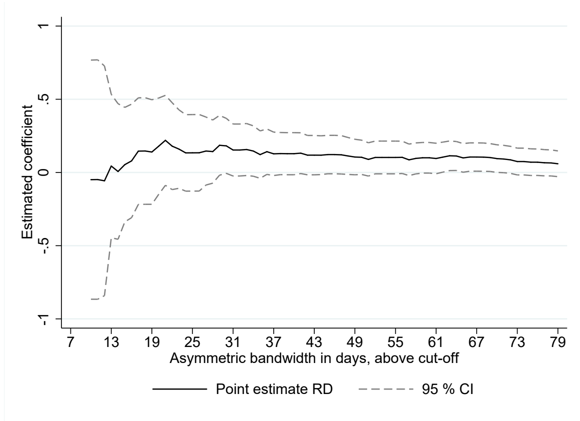

(b) With covariates, 2010

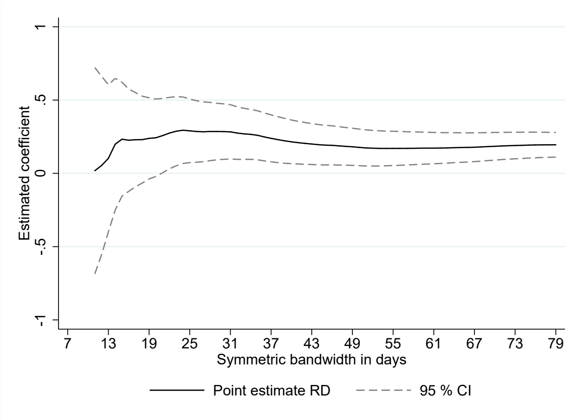

(c) Triangular kernel, 1994

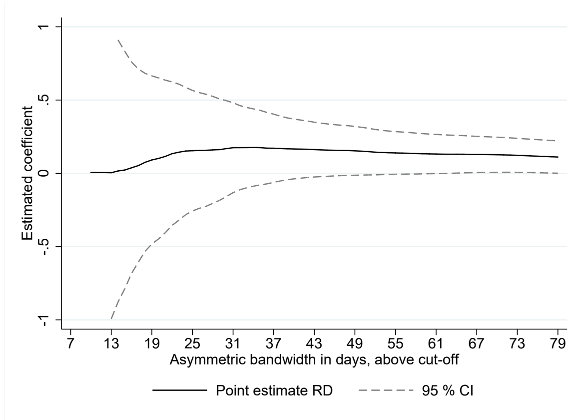

(d) Triangular kernel, 2010

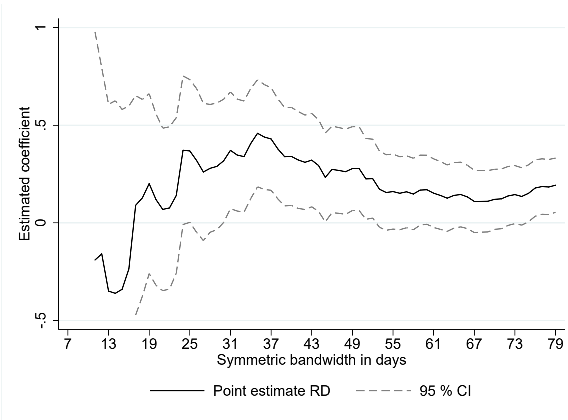

(e) Quadratic polynomial, 1994

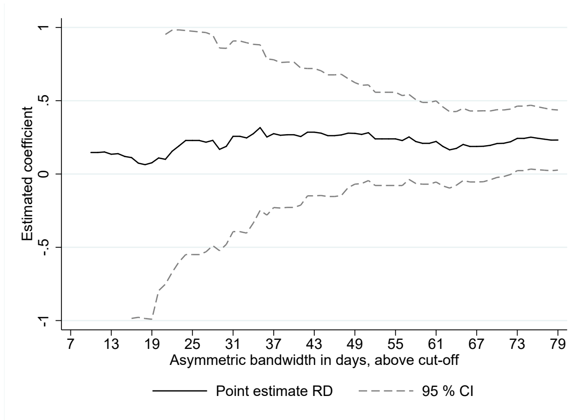

(f) Quadratic polynomial, 2010

Notes: Subfigures (a) and (b) display the main fuzzy RD estimates but with included covariates. Subfigures (c) and (d) employ a triangular kernel instead of a uniform one. Subfigures (e) and (f) employ a local quadratic polynomial instead of a local linear specification. For 1994, the bandwidth is symmetric, while it is asymmetric for 2010, with the upper limit always set at 30. CIs are capped at (-1,1)

Figure E4: Robustness analysis for small symmetric bandwidths below 30 days. 2010 sample

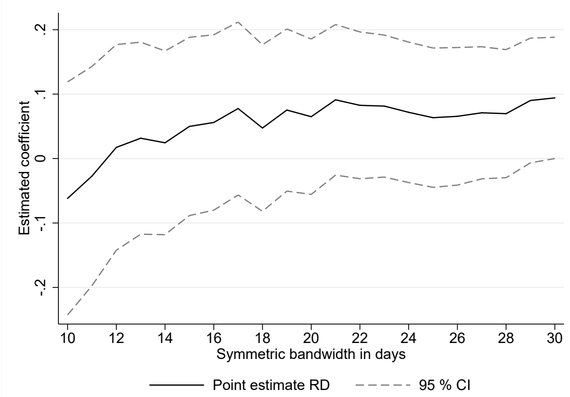

(a) Sharp estimates

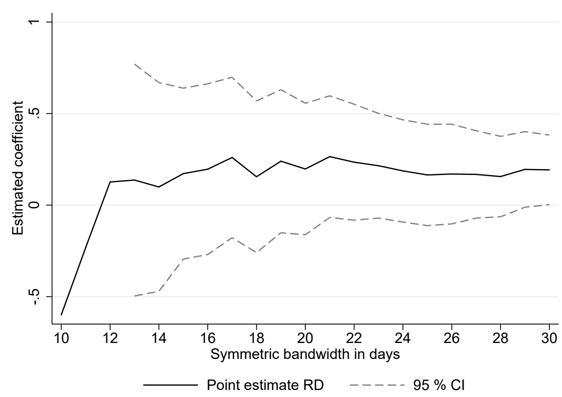

(b) Fuzzy estimates

Notes: The figure displays RD estimates for the 2010 sample with a symmetric bandwidth below 30 days. The RD estimate is the same as in the main analysis in Figure 4, where we do not include any covariates and use a uniform kernel. Note that the MSE-optimal bandwidth is above the maximum value of the running variable. CIs are capped at (-1,1)

Figure E5: “Honest” confidence intervals

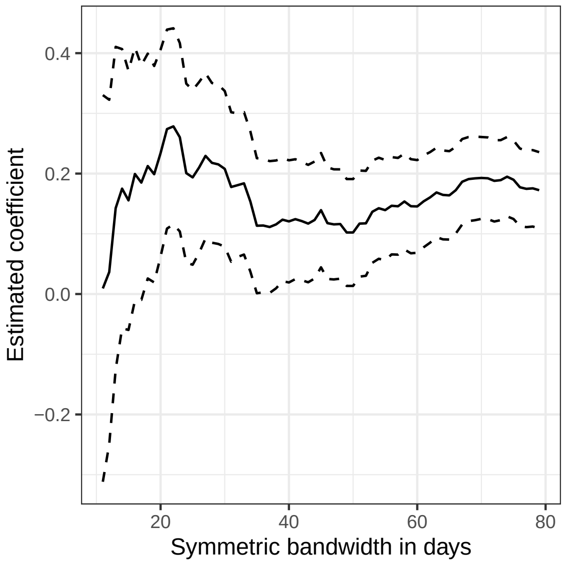

(a) Sharp, 1994

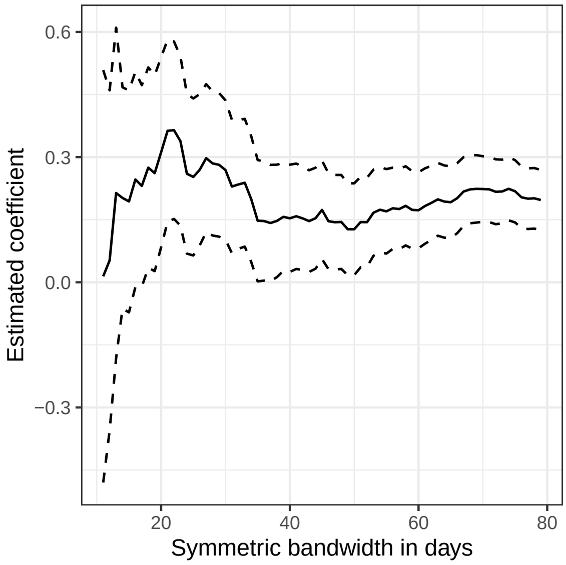

(b) Fuzzy, 1994

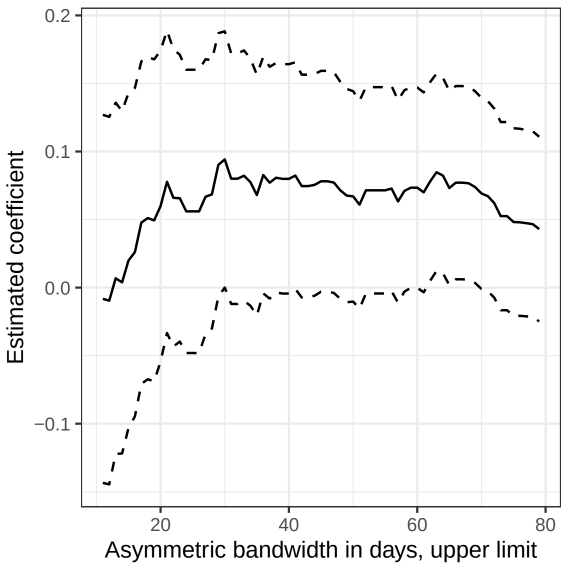

(c) Sharp, 2010

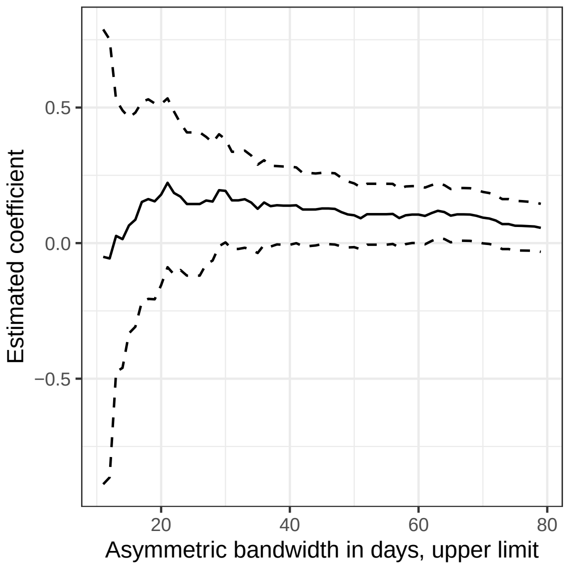

(d) Fuzzy, 2010

Notes: Honest confidence intervals according to Kolesár and Rothe (2018) and Armstrong and Kolesár (2020). The confidence intervals assume linearity with a bound of the second derivative at 0.

Figure E6: Bias-corrected robust confidence intervals

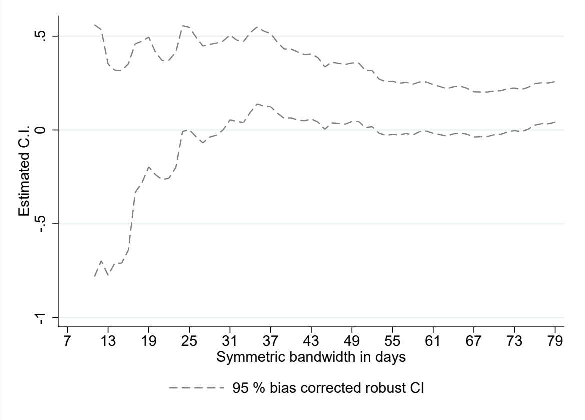

(a) Sharp, 1994

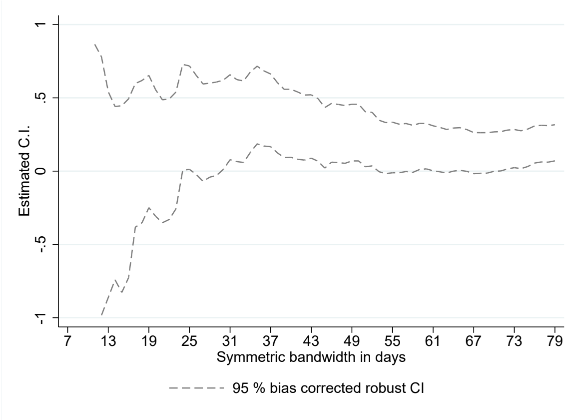

(b) Fuzzy, 1994

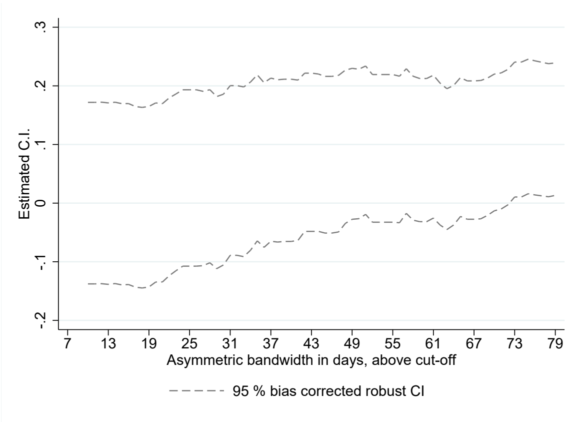

(c) Sharp, 2010

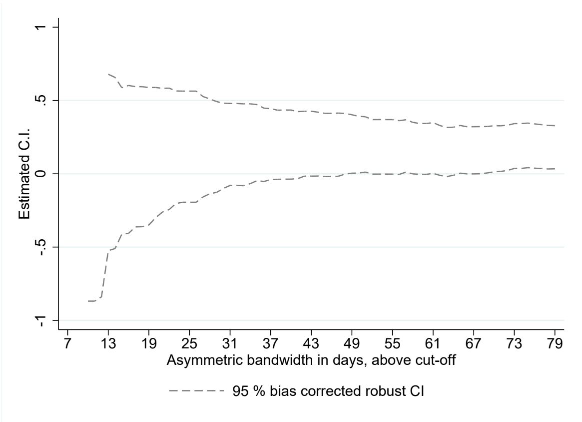

(d) Fuzzy, 2010

Notes: Bias-corrected robust confidence intervals according to Calonico et al. (2014b), Calonico et al. (2014a), and Calonico et al. (2017). Given that the authors of these papers argue for a separation between point estimation and inference and that the point estimates for the bias-corrected robust confidence intervals are not centered in the CIs for a given bandwidth, we only plot the CIs in the figure.

Figure E7: Robustness analysis excluding EU nationals

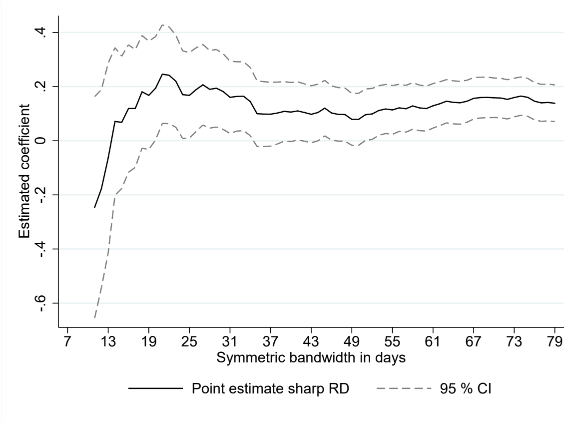

(a) Sharp, 1994

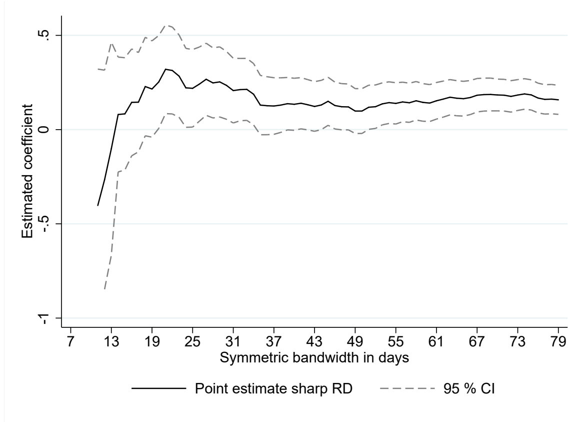

(b) Fuzzy, 1994

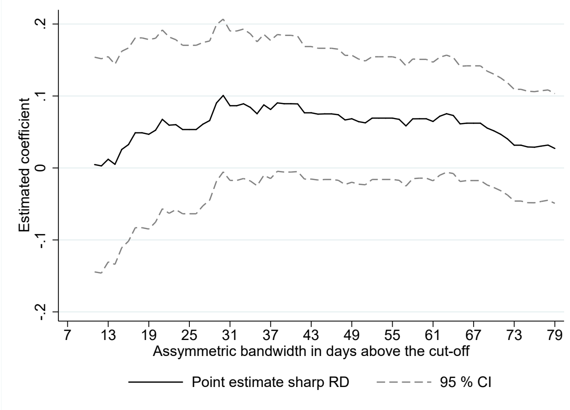

(c) Sharp, 2010

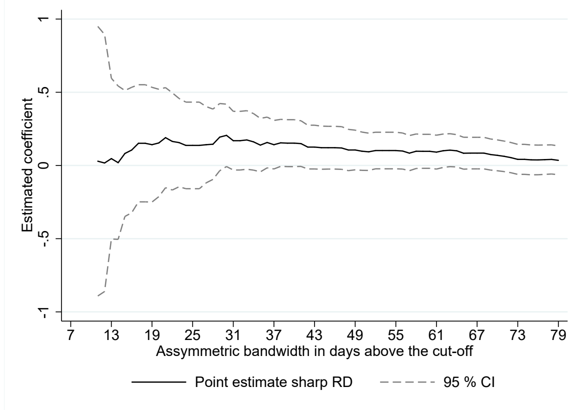

(d) Fuzzy, 2010

Notes: Robustness analysis excluding individuals born in another EU country. The RD estimate is otherwise the same as in the main analysis in Figure 4, where we do not include any covariates and use a uniform kernel. CIs are capped at (-1,1)

G Additional analyses and heterogeneity regarding the main results

New citizens in Sweden are able to vote in a more salient election, which, as we have proposed, subsequently increases turnout in the less salient elections. However, by definition, a new citizen also gets to vote in a particular election for the *first time*. An alternative hypothesis to the observed results is thus a "novelty effect," meaning that some are particularly prone to vote due to the novelty of partaking in national elections.

One way of evaluating the novelty hypothesis is to use information on date of citizenship in a prior election. We do this in Figure F1a by comparing turnout in the 1994 elections among those who received citizenship *i)* just before and *ii)* just after the elections in 1991. The latter of these groups will vote in national elections for the very first time in 1994. Hence, if there is a novelty effect, we expect those becoming citizens just after the elections in 1991 to vote *more* in 1994 than those who became citizens in time for the elections in 1991. However, Figure F1a does not suggest such a pattern. We estimate no effect at the cut-off and consequently find little support for a novelty story. Figure F1b shows similar estimates for turnout in 2010, using eligibility in the 2006 elections.

In Figure F2, we present additional disaggregated results for the analysis in the main text in Figure 5. Here in the Appendix, we split the analysis into non-naturalized EU citizens and non-naturalized non-EU citizens. As already discussed in the main text, voter turnout is particularly high in the 2003 referendum among non-naturalized EU citizens, for whom this topic is particularly salient.

The next natural question is whether the main effect is the same for all non-naturalized immigrants or whether a reform along these lines would instead mobilize particular subgroups among them. In an attempt to answer this question, we have performed some heterogeneity analyses. We examine heterogeneity with respect to five characteristics: age, socio-economic status (SES), region of birth, years since immigration, and the number of immigrants already in Sweden from a particular country (size of ethnic enclave).

In terms of age, we divide the sample into two groups: young (all below the age of

30) and old (those aged 30 and above). Socio-economic position is operationalized as whether the individual in question is employed (high SES) or not (low SES). In terms of region of birth, we split the sample into the same regions we used in the balance tests: Europe, Asia, Africa, and South America. Years since immigration is simply measured as the number of years since the immigrant first received a residence permit. We then split the sample in two: one group who became citizens in six years (the median value) or shorter and one group who became citizens after more than six years. Finally, we study the size of ethnic enclave by counting the number of individuals living in Sweden from a specific country group in 1994 and 2010. We rank the country groups in terms of size and split the countries into two groups .[13] The results are presented in Figures F3–F7. Given that this is a split-sample analysis with the same underlying total number of observations used in Figure 3, we go straight to the bandwidth graphs for this heterogeneity analysis. Please note that the CIs are capped at (-1,1) just as before in these figures.

The results are less precise here than for our main analysis. This is to be expected, given the few observations we now use. As a general takeaway, there is little evidence of heterogeneity. Neither the young, the non-employed, nor those receiving citizenship relatively quickly seem to be affected to any greater extent than the old, the employed, and citizens with a longer tenure before becoming citizens. Also, we see no statistical difference in effect size between those coming from countries with a large or small enclave living in Sweden at the time of voting.

When studying region of origin, we detect clear and positive jumps for all regions, with one exception: we detect no effect for individuals born in Africa. Admittedly, the heterogeneity analyses are performed using much smaller samples (this is particularly true for the sample of individuals born in Africa), and we need to be careful not to read in too much into the observed lack of positive effects for individuals born in Africa. Nevertheless, as we have documented no heterogeneity along age, socioeconomic status, time in country, or size of ethnic enclave, we believe that it is reasonable to rule out these mechanisms as

[13] In the minority of cases where immigrants have immigrated twice, we use the first recorded immigration date.

explanations. We can thus at least conclude that any explanation is likely beyond these measures.

Lastly, we run an analysis where we investigate spillover effects. Given that we have found evidence that eligibility in the more salient national election affects voter turnout in local elections, it is interesting to see whether this strong individual effect is transmitted to people in close proximity. If we find significant positive effects, this could mean that expanding voting rights has multiplier effects beyond the average effects found thus far. Second, assuming a positive spillover, the main effect in Figure 3 may be interpreted as an information effect, whereby non-naturalized immigrants do not vote due to a lack of information. The main effect may in that case partly be interpreted as information acquisition at the time of gaining the right to vote in national elections. If that is the case, it is likely that the effect should be transmitted to others in close proximity.

We analyze this by replacing the outcome variable to voter turnout in local elections among 1) spouses that became citizens in the same year and 2) immigrants from the same immigrant group who live in the same parish (as the individual in question) who also became citizens in 1994 and 2010, respectively. The results are presented in Figure F9. We find no clear evidence of such spillover effects.[14]

Figure F1: Novelty hypothesis

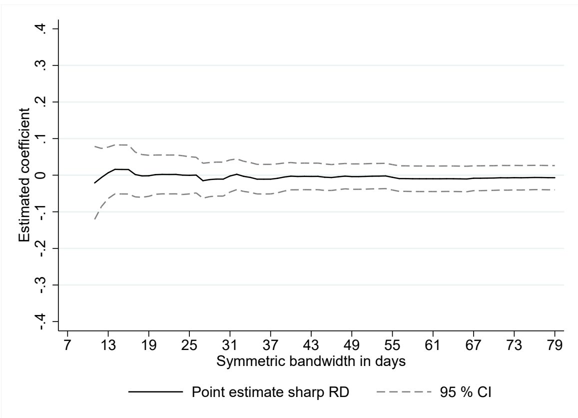

(a) Threshold 1991, outcome 1994

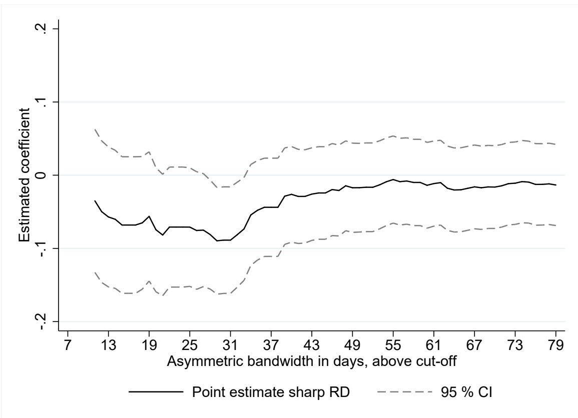

(b) Threshold 2006, outcome 2010

Notes: The figure displays RD estimates using turnout in the 1994 (a) and 2010 (b) elections. The running variable is defined using citizenship acquisition in connection with the prior election. The RD plot is otherwise the same as in the main analysis in Figure 4.

[14] Corresponding bandwidth graphs with regression results are available upon request.

Figure F2: Turnout in 2002 local elections and 2003 euro referendum

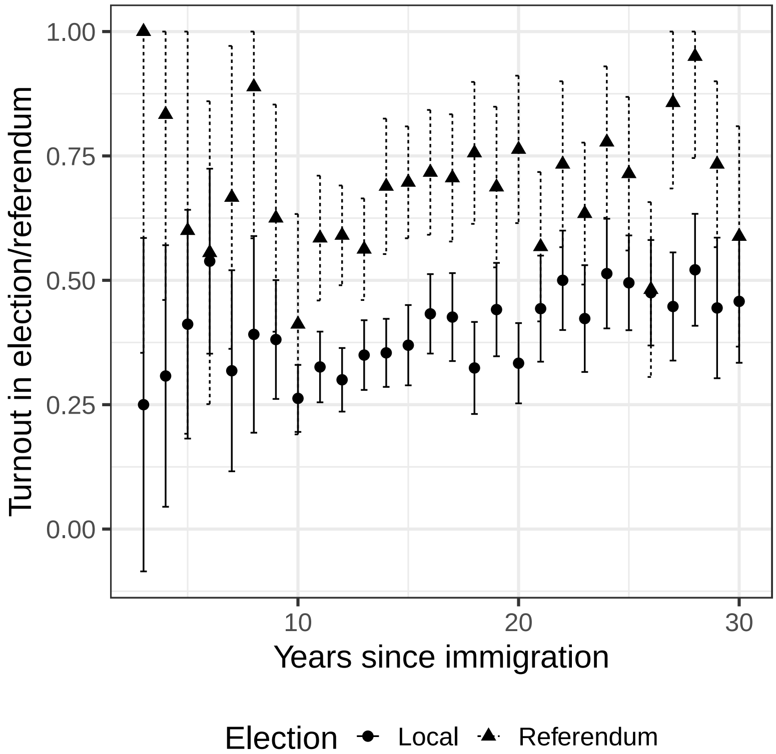

(a) Non-naturalized EU citizens

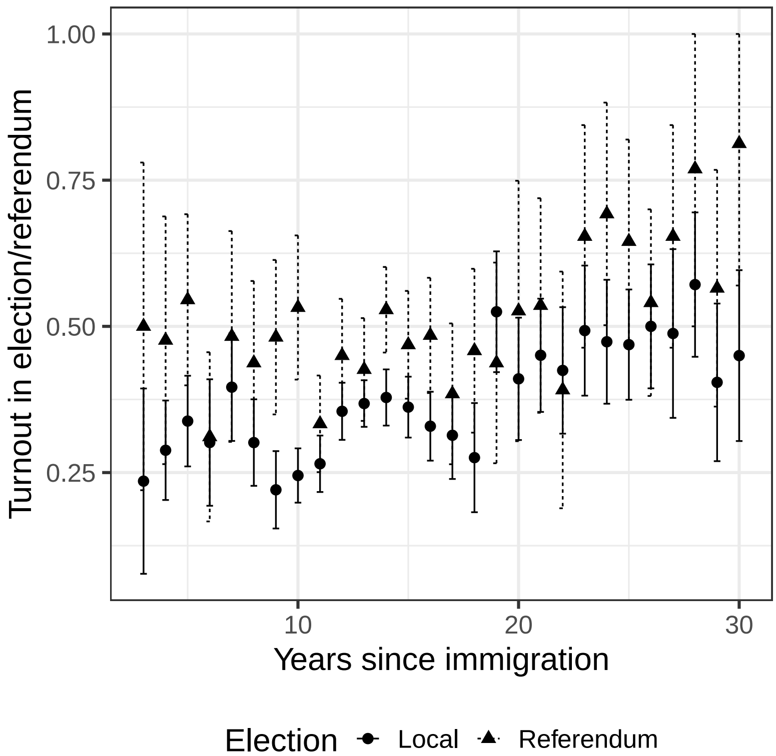

(b) Non-naturalized non-EU citizens

Notes: Turnout among non-naturalized immigrants in the 2002 local elections and 2003 euro referendum.

Figure F3: Heterogeneity analysis. Sharp estimates. Age and SES

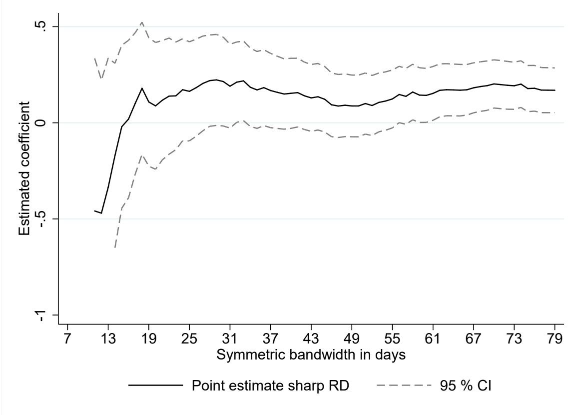

(a) Below 30, 1994

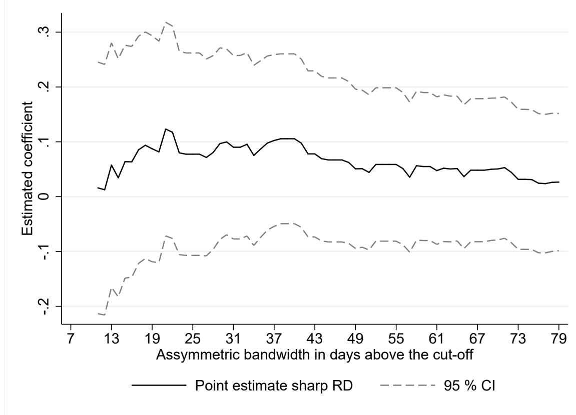

(b) Below 30, 2010

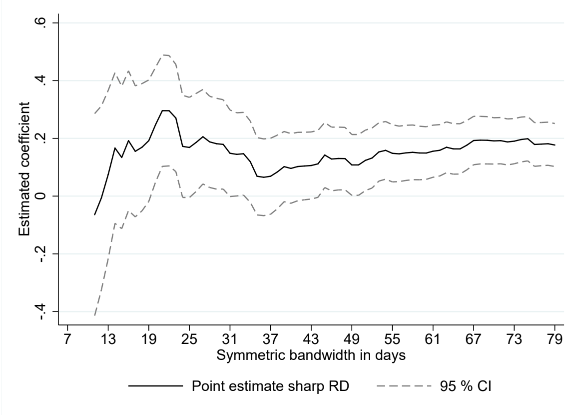

(c) Above 30, 1994

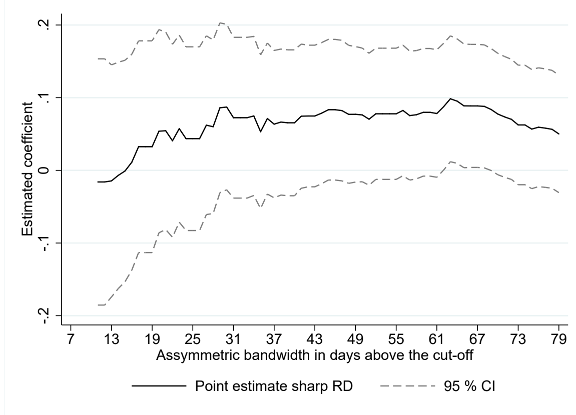

(d) Above 30, 2010

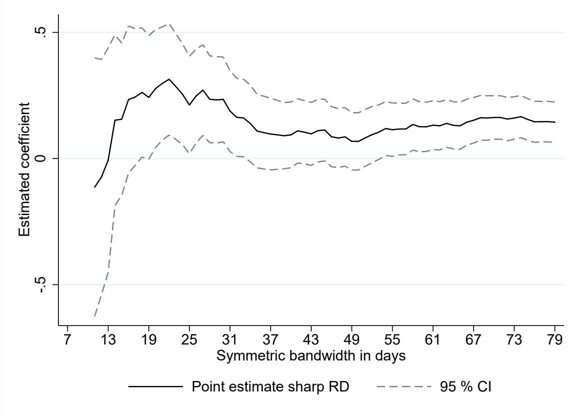

(e) Unemployed, 1994

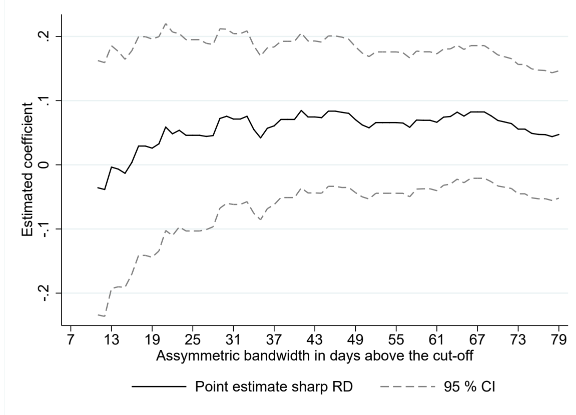

(f) Unemployed, 2010

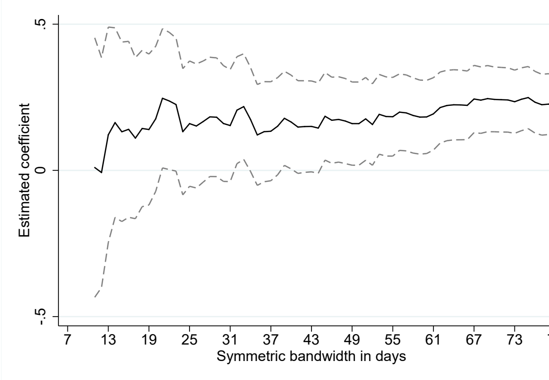

(g) Employed, 1994

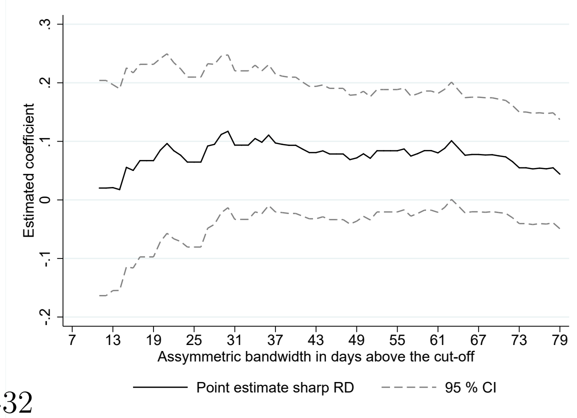

(h) Employed, 2010

Figure F4: Heterogeneity analysis. Fuzzy estimates. Age and SES

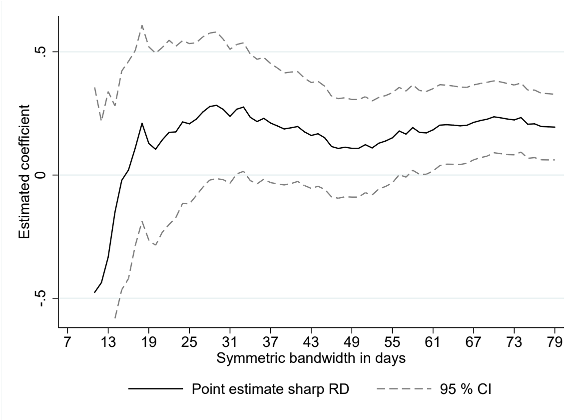

(a) Below 30, 1994

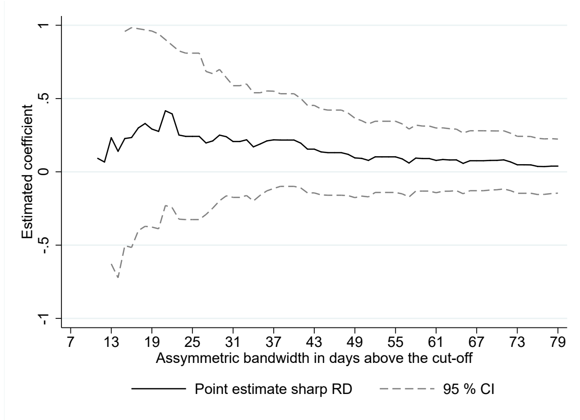

(b) Below 30, 2010

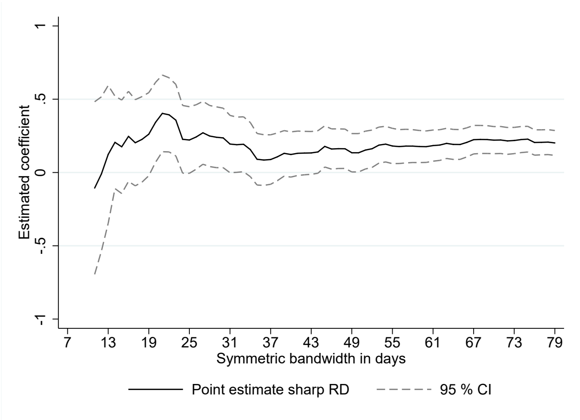

(c) Above 30, 1994

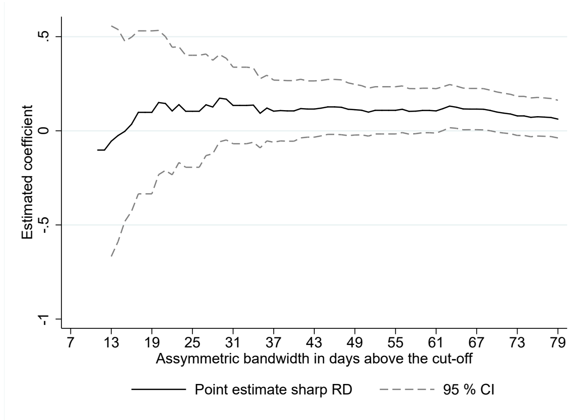

(d) Above 30, 2010

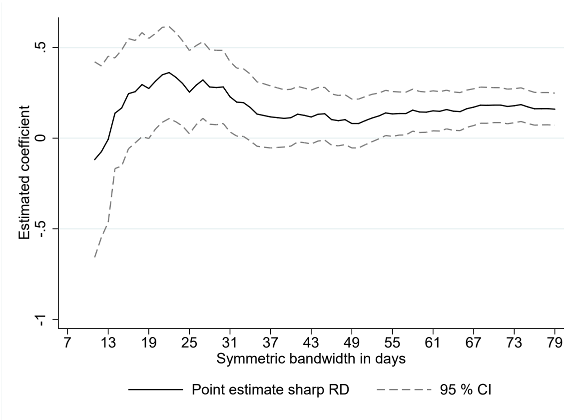

(e) Unemployed, 1994

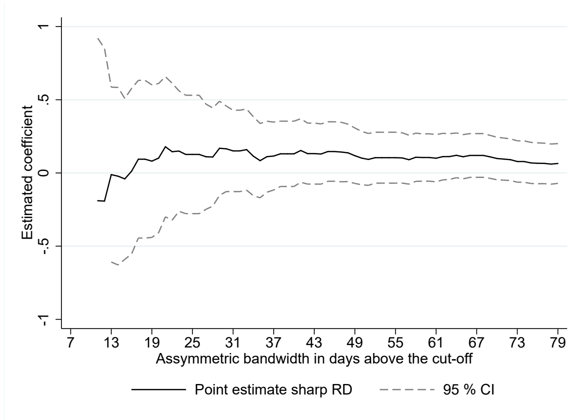

(f) Unemployed, 2010

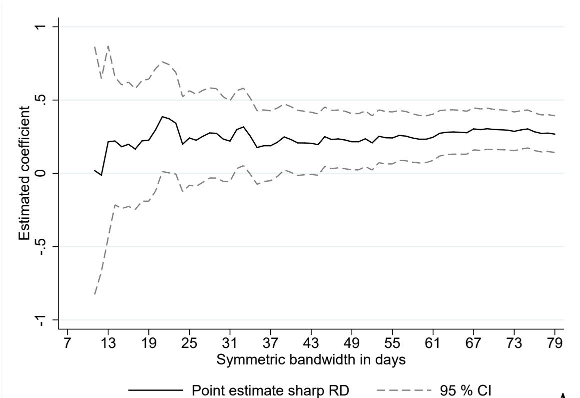

(g) Employed, 1994

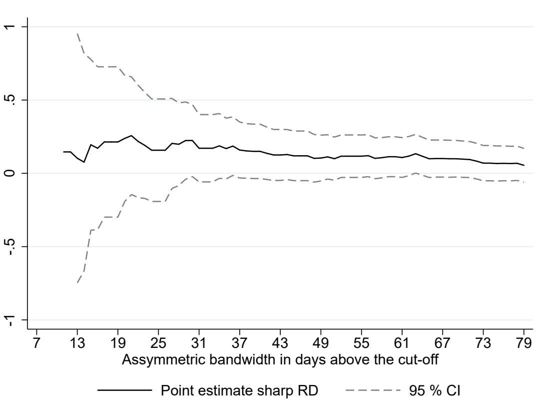

(h) Employed, 2010

Figure F5: Heterogeneity analysis. Sharp estimates. Region of origin

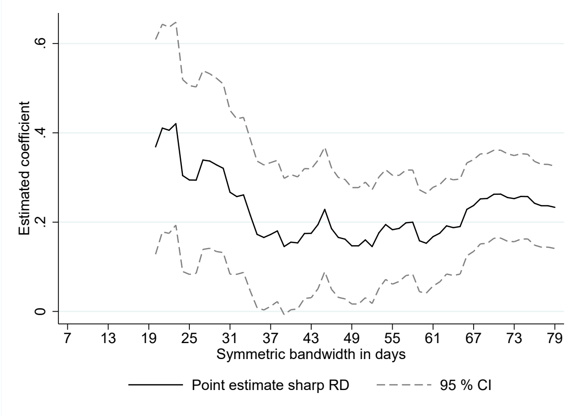

(a) Europe, 1994

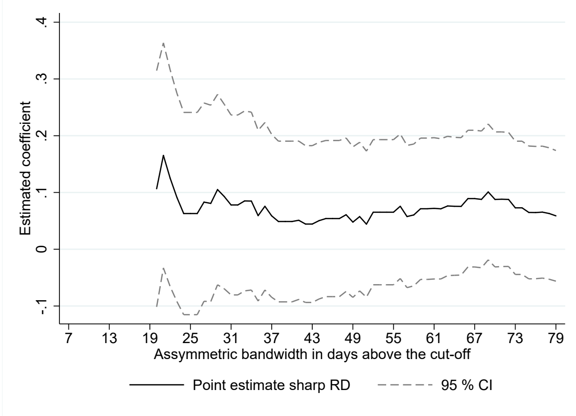

(b) Europe, 2010

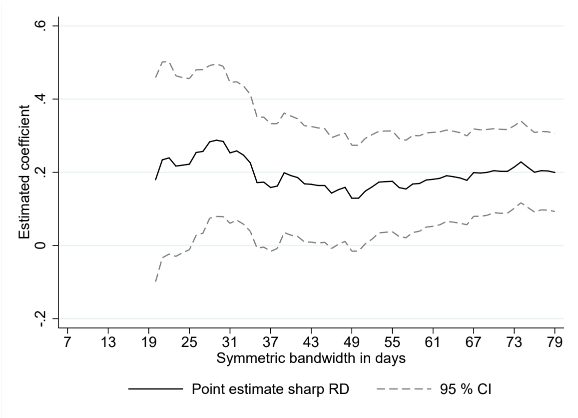

(c) Asia, 1994

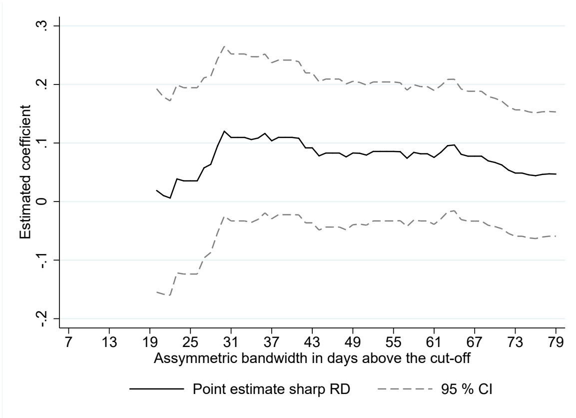

(d) Asia, 2010

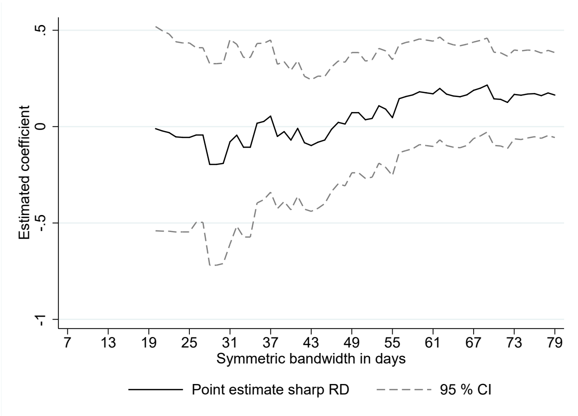

(e) South America, 1994

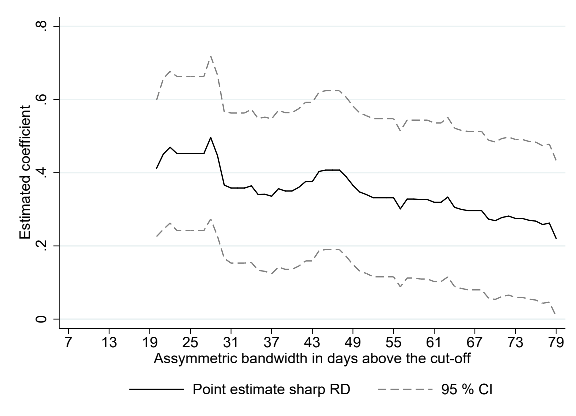

(f) South America, 2010

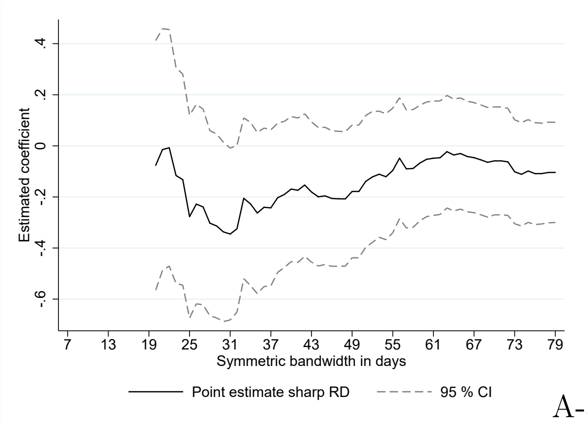

(g) Africa, 1994

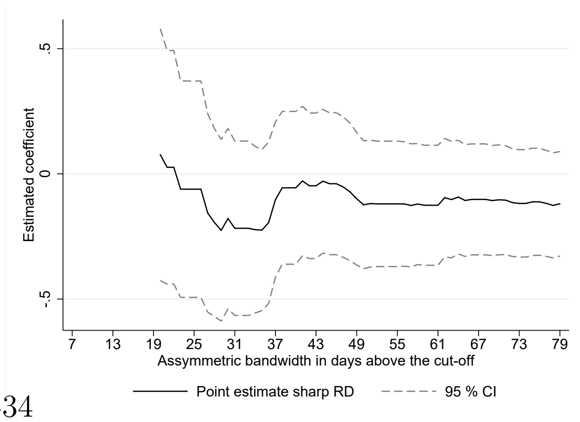

(h) Africa, 2010

Figure F6: Heterogeneity analysis. Fuzzy estimates, region of origin

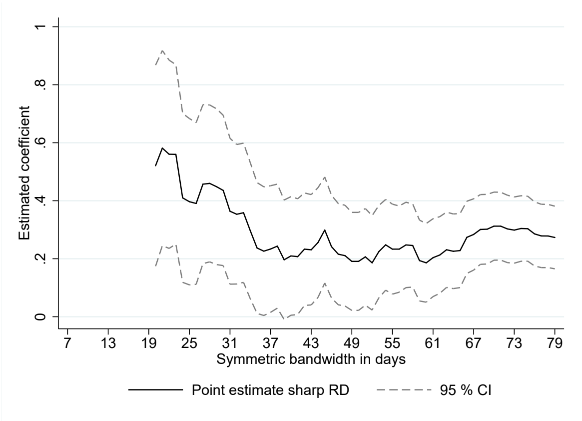

(a) Europe, 1994

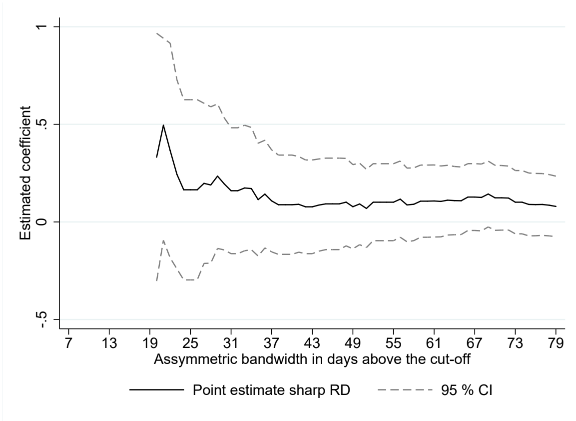

(b) Europe, 2010

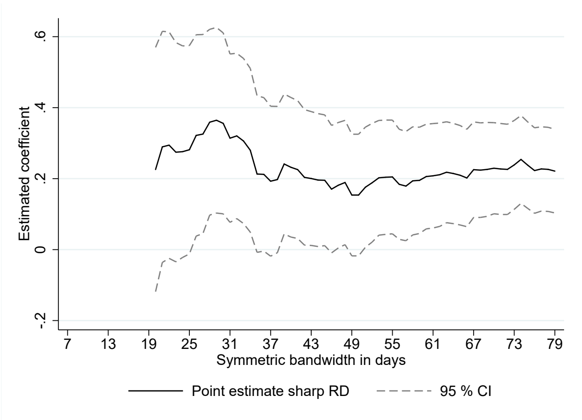

(c) Asia, 1994

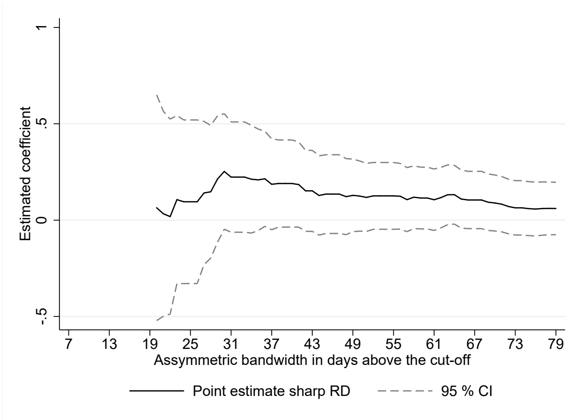

(d) Asia, 2010

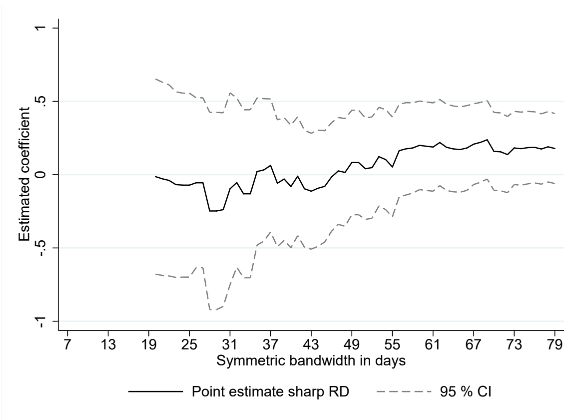

(e) South America, 1994

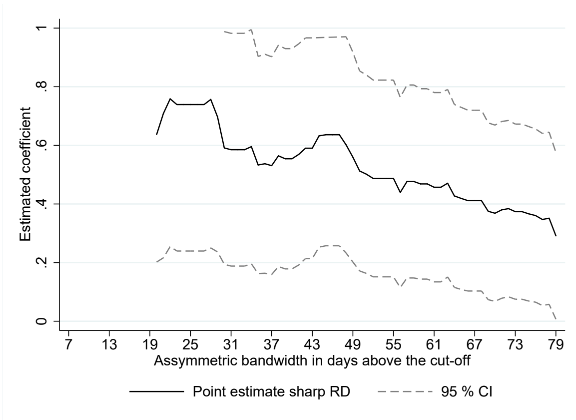

(f) South America, 2010

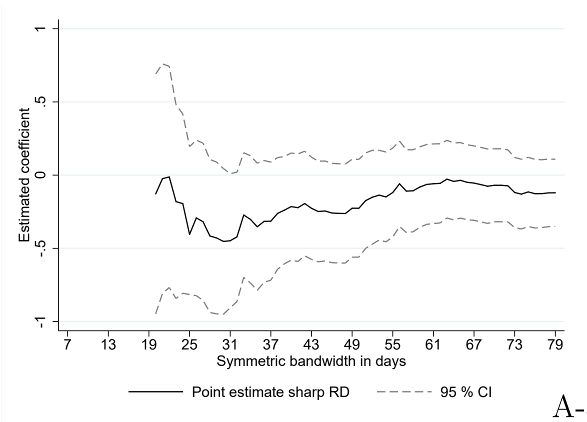

(g) Africa, 1994

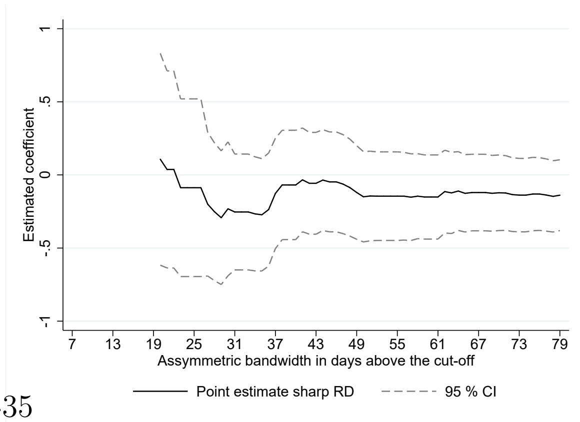

(h) Africa, 2010

Figure F7: Heterogeneity analysis. Time in Country

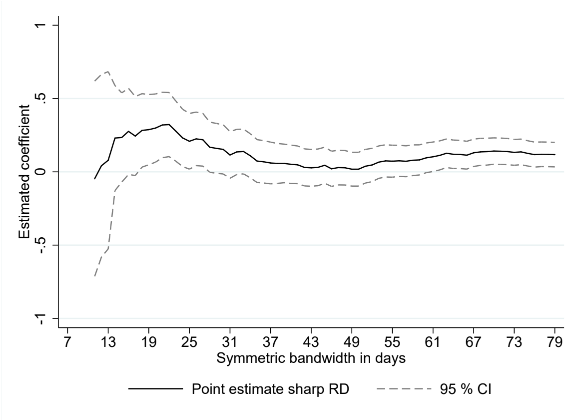

(a) Citizen in ≤ 6 years, 1994, sharp

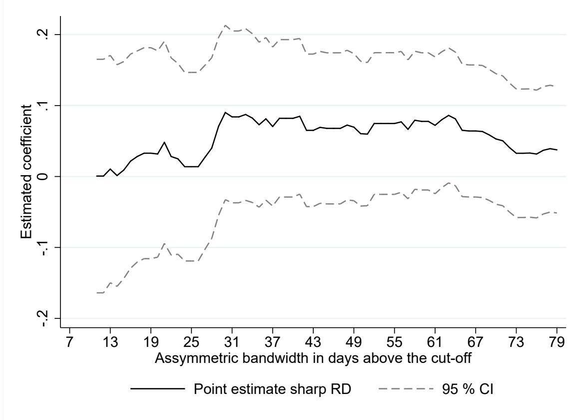

(b) Citizen in ≤ 6 years, 2010, sharp

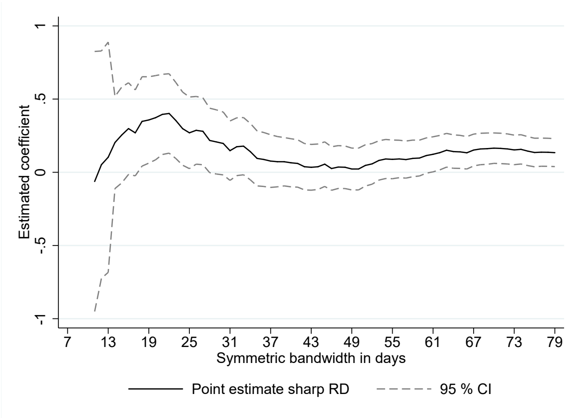

(c) Citizen in ≤ 6 years, 1994, fuzzy

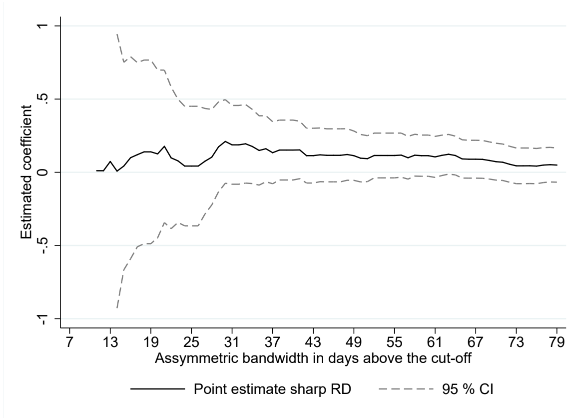

(d) Citizen in ≤ 6 years, 2010, fuzzy

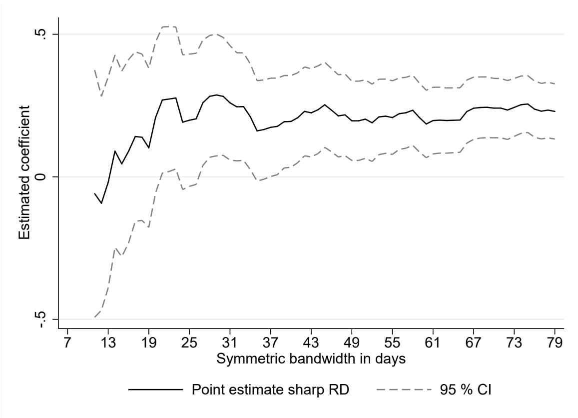

(e) Citizen in > 6 years, 1994, sharp

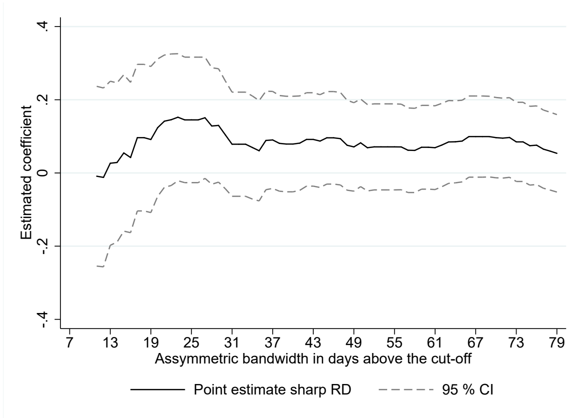

(f) Citizen in > 6 years, 2010, sharp

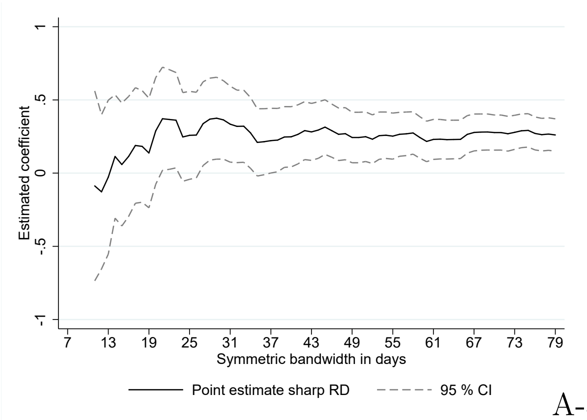

(g) Citizen in > 6 years, 1994, fuzzy

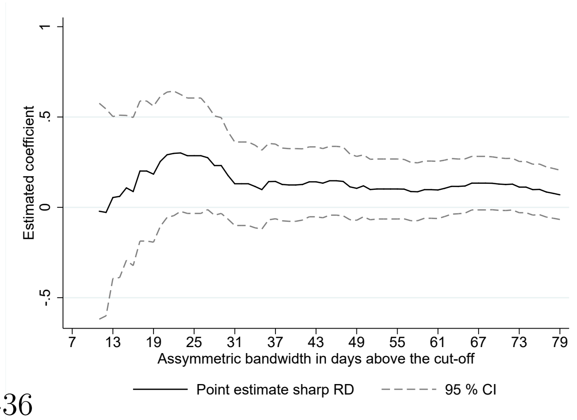

(h) Citizen in > 6 years, 2010, fuzzy

Figure F8: Heterogeneity analysis. Large and small enclave

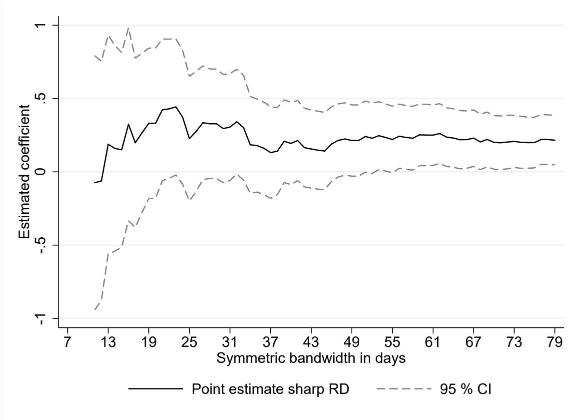

(a) Small enclave, 1994, sharp

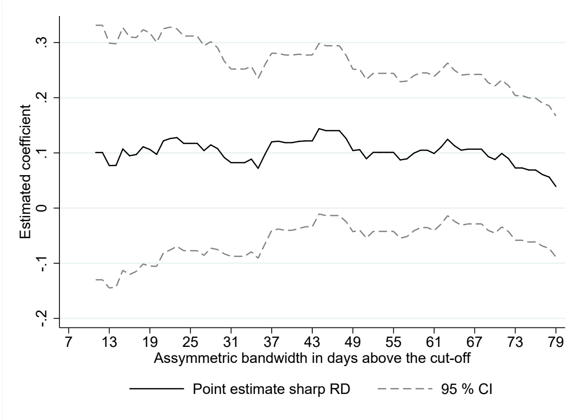

(b) Small enclave, 2010, sharp

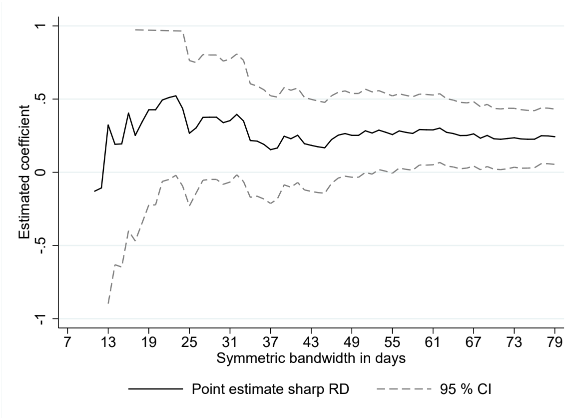

(c) Small enclave, 1994, fuzzy

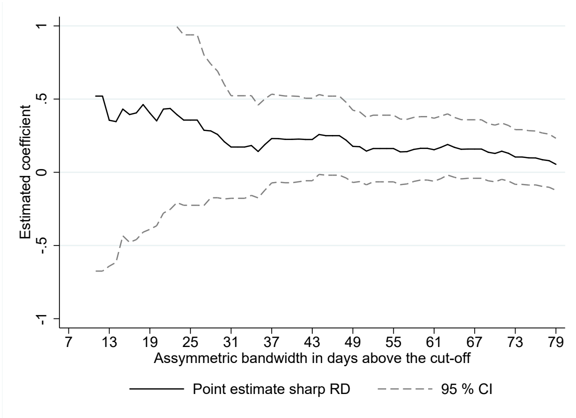

(d) Small enclave, 2010, fuzzy

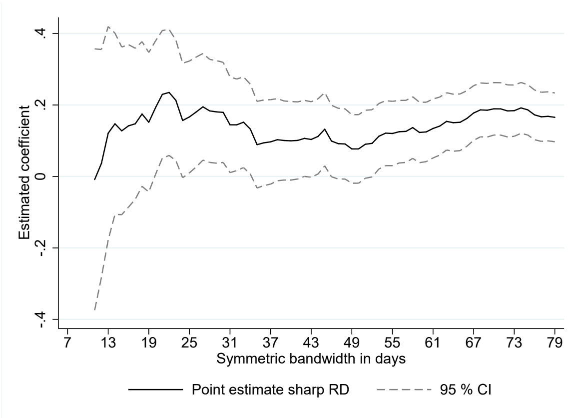

(e) Large enclave, 1994, sharp

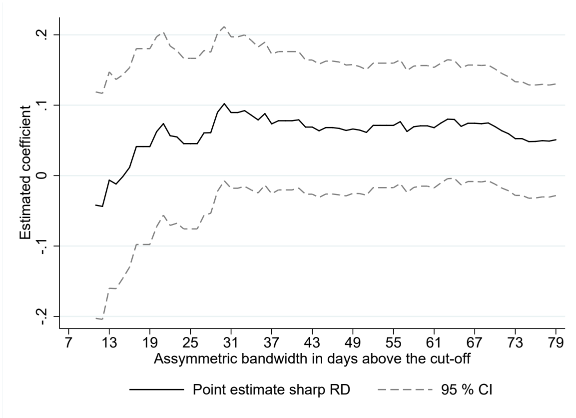

(f) Large enclave, 2010, sharp

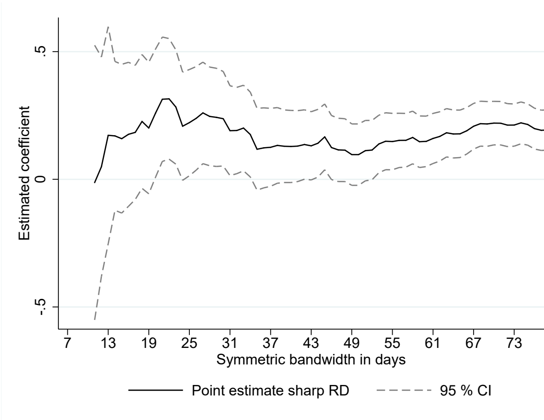

(g) Large enclave, 1994, fuzzy

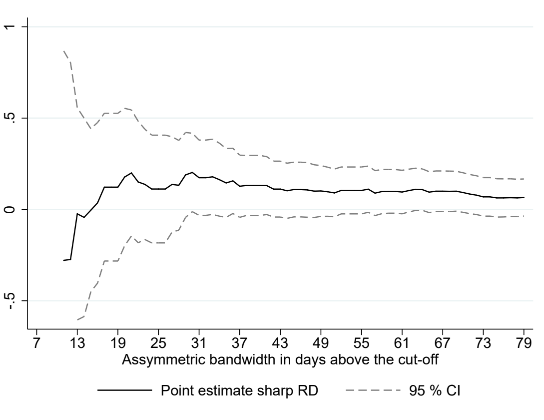

(h) Large enclave, 2010, fuzzy

Figure F9: Spillover effects on spouses and other immigrants

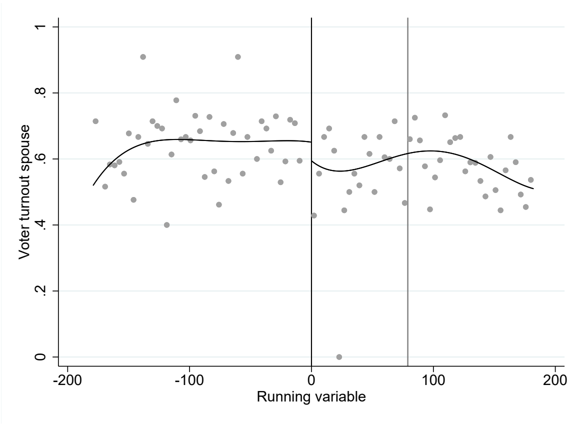

(a) 1994 municipal elections

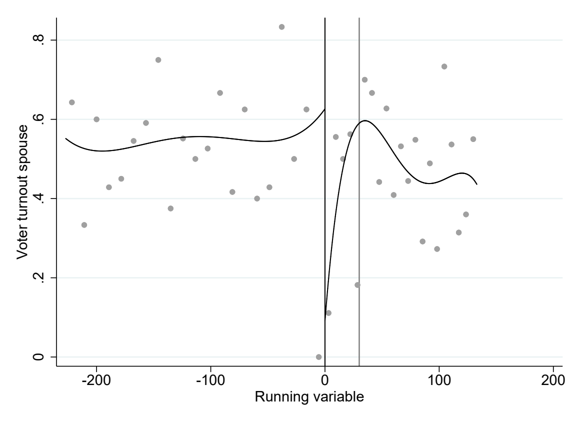

(b) 2010 municipal elections

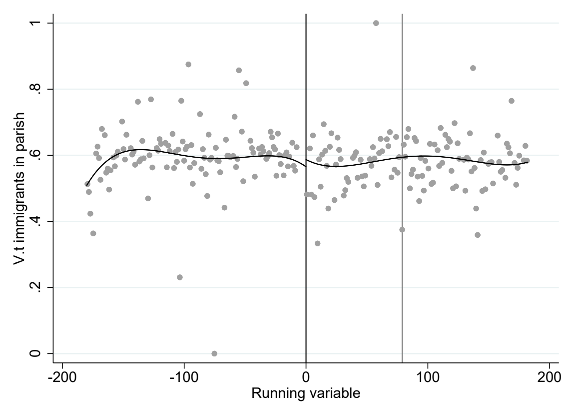

(c) 1994 municipal elections

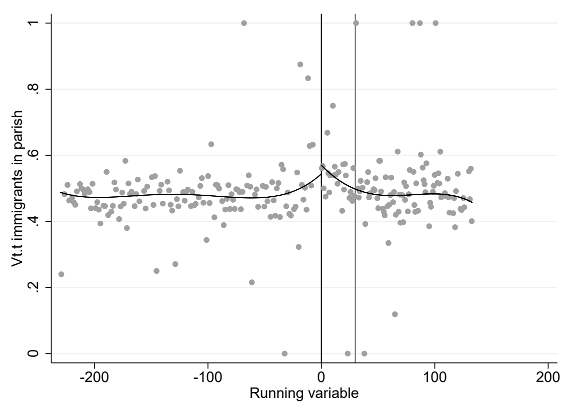

(d) 2010 municipal elections

Notes: RD plots using observations from the 1994 (a) and 2010 (b) elections. The running variable measures days before/after the eligibility cut-off.

H Analyzing the pure citizen effect

On several occasions by now, we have noted that becoming a citizen per se and gaining the right to vote are two different treatments that often go hand in hand. One of the strengths of our article is that we are able to disentangle these two elements from each other since we have access to individuals who became citizens too close to the election to be registered in the electoral rolls.

In Figure G1, we focus on individuals who gained their citizenship prior to the election but too close for it to be registered in the electoral rolls and compare with those who gained their citizenship after the election. We run these regressions using the same donut specification as in our main analysis in Figure 3. One difference is that we do not run a fuzzy specification. The reason for this is that we do not have any variable that we can use, since our only observation concerns the date of citizenship. The calculated optimal bandwidths for the pure citizenship treatment are 21 days for the entire sample, 24 days for 1994, and 25 days for 2010, respectively.

Our conclusion from Figure G1 is that there seems to be no pure citizenship effect.

Figure G1: The pure citizenship effect

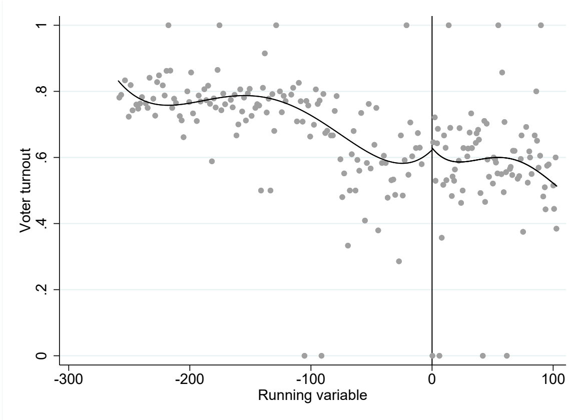

(a) 1994, RD plot

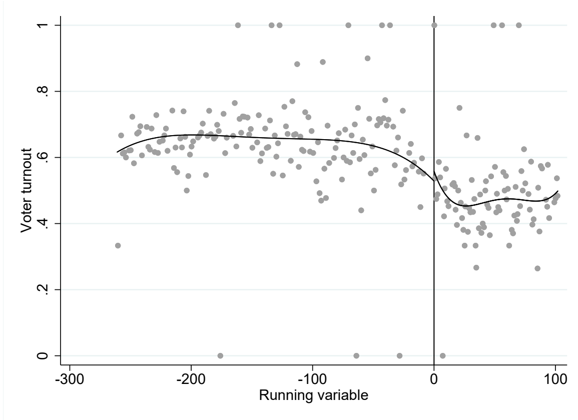

(b) 2010, RD plot

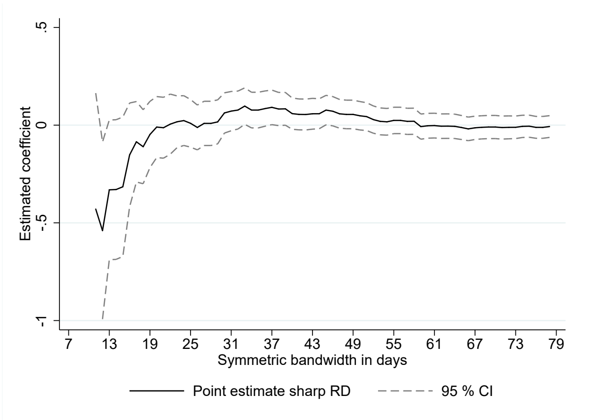

(c) 1994, bandwidth graph

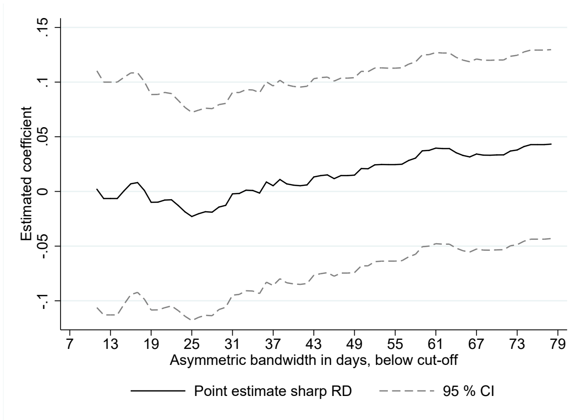

(d) 2010, bandwidth graph

Notes: The RD estimation is the same as in the main analysis but for a different treatment (pure citizenship effect), where we do not include any covariates and use a uniform kernel. In the bottom two figures, the CIs are capped at (-1,1)

References

Armstrong, Timothy B and Michal Kolesár. 2020. "Simple and honest confidence intervals in nonparametric regression." *Quantitative Economics* 11 (1): 1–39.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik. 2014a. "Robust data-driven inference in the regression-discontinuity design." *The Stata Journal* 14 (4): 909–946.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014b. "Robust nonparametric confidence intervals for regression-discontinuity designs." *Econometrica* 82 (6): 2295–2326.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2015. "Optimal data-driven regression discontinuity plots." *Journal of the American Statistical Association* 110 (512): 1753–1769.

Calonico, Sebastian, Matias D Cattaneo, Max H Farrell, and Rocio Titiunik. 2017. "rdrobust: Software for regression-discontinuity designs." *The Stata Journal* 17 (2): 372–404.

Dehdari, Sirus H., Jaakko Meriläinen, and Sven Oskarsson. 2021. "Selective abstention in simultaneous elections: Understanding the turnout gap." *Electoral Studies* 71, doi.org/10.1016/j.electstud.2021.102302.

European Council. 1994. *Document 31994L0080: Council Directive* 94/80/*EC of 19 December 1994*. https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:31994L0080. Accessed: 2020-08-31.

Gelman, Andrew and Guido Imbens. 2019. "Why high-order polynomials should not be used in regression discontinuity designs." *Journal of Business & Economic Statistics* 37 (3): 447–456.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw. 2001. "Identification and estimation of treatment effects with a regression-discontinuity design." *Econometrica* 69 (1): 201–209.

Kolesár, Michal and Christoph Rothe. 2018. "Inference in regression discontinuity designs with a discrete running variable." *American Economic Review* 108 (8): 2277–2304.

Lee, David S and David Card. 2008. "Regression discontinuity inference with specification error." *Journal of Econometrics* 142 (2): 655–674.

McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: A density test." *Journal of Econometrics* 142 (2): 698–714.