

SI1 Checking for household relationships

This project assumes that voters and criminal defendants living at the same geographic coordinates (according to public records) know each other, and will often have some sort of familial or romantic relationship. I check this by drawing a random sample of observations from my dataset and verifying that many of the households do appear to include such relationships. I run this check for households of 2009-2012 defendants. I randomly sample 100 rows from the sample. Then, for each row, I investigate the relationship between the registered voter and the criminal defendant who live at this address. First, I check whether they share a last name; this is a fairly conservative measure of family relationships, as many people may be related but not share a surname (and the probability of coincidences in which people report the same residential address and share a last name but are not related seems low). 47% of households contain voters and defendants that share a last name. This is quite high, considering that many romantic partners in this sample may be unmarried. Next, for households that do not share a last name, I look for other evidence of connections in public records of marriages and births. By looking up defendants' and voters' names in the Texas state birth index from 1907-1993 (searchable through Familysearch.org), I find evidence of parent-child relationships, shared children, or shared parents (sibling relationships) for an additional 10% of observations. A number of other observations appear to be related (age gaps suggest parent-child or parent- grandchild relationships, and naming similarities suggest a familial link), but could not be verified using public records and are not counted here. In total, I find strong evidence of close familial relationships for 57% of observations in this random sample. This is quite high, as it does not capture people cohabiting without children in common, and may also not capture parent-child relationships for people born outside Texas, not to mention close friends or other relatives.

SI2 About the Sample

SI2.1 Geocoding Voters and Defendants

The same process was used to assign latitude/longitude to both registered voters in Harris County (using their addresses from the Texas voter file) and people facing criminal cases in Harris County (using their addresses from court records). In both cases, the basic geocoding procedure was:

- Preprocess geographic data in R to remove rows with blank addresses and drop extraneous columns
- Load a comma-separated file of addresses into ArcGIS 10.4.1, and geocode them using an address locator provided by Harvard Center for Geographic Analysis and a match threshold of 49
- Merge the geocoded addresses back into the main dataset

Once both datasets had geographic coordinates, defendants were matched to registered voters based on geographic proximity: anyone living at the same location (within 5 meters) of each other were considered to be household members. A few types of apparent matches were then discarded: first, locations with more than 10 registered voters living at the same point were omitted, because of concerns that these places represent large apartment buildings or complexes where people may not be living in the same household. And second, defendants who were themselves registered voters and so were “matched” to themselves were removed from the dataset (using name and birthdate), because the treatment of interest is proximal contact, not personal contact with the criminal legal system.

SI2.2 Sample sizes

The analysis presented here rely on data for a set of people who 1) are registered voters in Harris County, 2) have reliable street addresses on their voter records that can be geocoded, and 3) live at the same address as someone who was charged with their first misdemeanor crime in 2012-2013 (and who had this address listed in their court record). The focus on registered voters misses some other household members with proximal contact, though (as I discuss further in Section 4 of the main manuscript), registered voters seem like a reasonable first place to look for demobilizing effects. But it is also worth thinking through how many people meet all three of the criteria above, and where people drop out of the sample.

The number of people who fall into these categories depends on how broad a time window we use: if focusing on only people whose household members' cases take place in the month around Election Day, we only observe about 1550 registered voters with this experience (as seen in Table 2). But if looking at all 2012-2013 misdemeanor cases, we can match 19,192 registered voters to a first-time misdemeanor defendant, a fairly substantial number.

Another way to think about who ends up in this sample is to start with defendants. Of the 49,179 people who faced misdemeanor charges for the first time in 2012 or 2013, 32,974 of them were successfully geocoded to an address in Harris County. Defendants that could not be geocoded were either listed as homeless, had a valid address that fell outside the county, had no address on file, or had (in relatively few cases) had an extremely vague or malformed address on file. Though it is possible that some defendants do genuinely live in households in Harris County but are not being geocoded (due to missing or incomplete addresses), it appears that most non-geocoded defendants either live outside the county or are homeless.¹ Of these 32,974 geocoded defendants, 9,673 of them were matched to a household with at least one registered voter. This is a fairly low match rate, but not an

¹It is possible that people listed as “homeless” still have family relationships with, and occasionally stay with, registered voters in the county, but it seems reasonable to imagine that these household ties are different (and might yield different/smaller demobilizing effects) than regularly living in the same home as someone.

implausible one: some defendants may live alone, while many others likely live in households where no one is registered to vote. One notable source of lost matches, though, is the omission of addresses with many registered voters, noted in Section 2 of the manuscript. The main merge specification used here drops any addresses with more than 10 registered voters, because of the difficulty of determining whether matched people actually live in the same household, or simply in the same apartment building. The merge is not especially sensitive to the exact cutpoint chosen here: dropping it to 8 people, or raising it to 12 people at an address, only changes the size of the resulting matched dataset by 4 or 7 percent, respectively. Even doubling the cutoff, to 20 registered voters at the same address, does not massively increase the number of matches (the resulting dataset is 14% bigger than the one used in the main analyses).

SI2.3 Generalizing Beyond Harris County

One concern about the estimates presented in the main paper is that they may not generalize beyond Texas. Harris county is unrepresentative of the US on a number of dimensions: it is a very large county, and a racially and ethnically diverse one, with a notoriously dangerous jail system and relatively low voter turnout. How should we think about these differences and what they mean for generalizability?

It is worth thinking through the direction of the “bias” that these characteristics could introduce to the estimates: do any of them introduce reasons to expect that the mainly-null results presented here would actually be much larger in other places? For most of these characteristics, the answer appears to be no. For example, one possible concern is that the “treatment” (of, for example, having a household member incarcerated) received in Harris County might differ from other places. If the treatment were harsher elsewhere, we might expect that the null results here would not replicate and that that stronger treatment would yield larger and longer-lasting treatment effects. But if anything, it seems likely that the

“treatment” would be harsher in Harris County than elsewhere. The conditions under which defendants sentenced to jail (and those that could not make bail) were held during this period were harsh, and seem likely to have been alarming both to defendants and to their families. For example, a 2010 letter from the Department of Justice Civil Rights Division to county officials reported:

“...we also conclude that certain conditions at the Jail violate the constitutional rights of detainees. Indeed, the number of inmates [sic] deaths related to inadequate medical care, described below, is alarming. As detailed below, we find that the Jail fails to provide detainees with adequate: (1) medical care; (2) mental health care; (3) protection from serious physical harm; and (4) protection from life safety hazards.”²

It is also worth noting that this dataset was collected before recent progressive changes to the Harris County courts’ handling of misdemeanor cases, such as the recent court ruling that holding misdemeanor defendants on cash bail (which would often be raised by their family members, or would result in defendants being detained pretrial) was unconstitutional.³

Similarly, it does not seem especially likely that the racial diversity of the county is somehow generating null results that would not be observed in a county with, for example, more white residents. If the literature yields any theoretical prediction about whether the effects of proximal contact differ by race, it is that these effects should be *larger* for people of color, and for African-Americans in particular⁴ Further, Section SI7 of this SI examines heterogeneity by race in the present data, and finds no evidence of different patterns for white and black voters.

²The full text of the letter is available from: https://www.justice.gov/sites/default/files/crt/legacy/2010/12/15/harris_county_jail_findlet_060409.pdf

³For more information on this case, see: <https://www.chron.com/news/houston-texas/houston/article/Harris-County-bail-system-unconstitutional-11108210.php>, <https://www.nytimes.com/2017/04/29/us/judge-strikes-down-harris-county-bail-system.html>

⁴See, for example, Maltby (2017) or White (Forthcoming).

Perhaps the most worrying dimension on which Harris County differs from other parts of the United States is voter turnout: about 39% of the voting-age population of the county turned out in 2016, compared to about 55% nationwide.⁵ It is plausible that places with higher baseline voter turnout would be more likely to show large or lasting demobilization effects from proximal contact. It is difficult to evaluate just how different Harris County is from other jurisdictions on this dimension, because what we'd really like to know is not overall turnout, but turnout rates among the (registered) household members of people facing criminal cases in other places. There is no reliable national data available on voting by this population, making it hard to evaluate this problem.

Still, it bears noting that prior voting turnout among the household members included in this study is low (as noted in Table 1 of the paper, 49% of them voted in 2008, compared to 55% for all registered voters in the county), but not anywhere near zero. It does not seem that these voters should be experiencing some sort of “floor effect” where there is no room for them to be further demobilized. Indeed, part of the rationale for selecting already-registered voters for the sample was that this was a population where *ex ante*, we would probably expect a relatively large effect.

Finally, Section SI7 of this SI looks for heterogeneity across past voting behavior, and does not find larger or longer-lasting effects among people who voted in 2008. If there were a large difference between regular voters and the rest of the sample, we ought to worry about whether higher-turnout jurisdictions would show dramatically different effects from Harris County. But without such a difference, these concerns seem less pressing.

⁵Note that these numbers look substantially lower than those reported in the paper because of the use of voting age population (VAP) as the denominator here, rather than registered voters.

SI3 Balance on Prior Voter Turnout

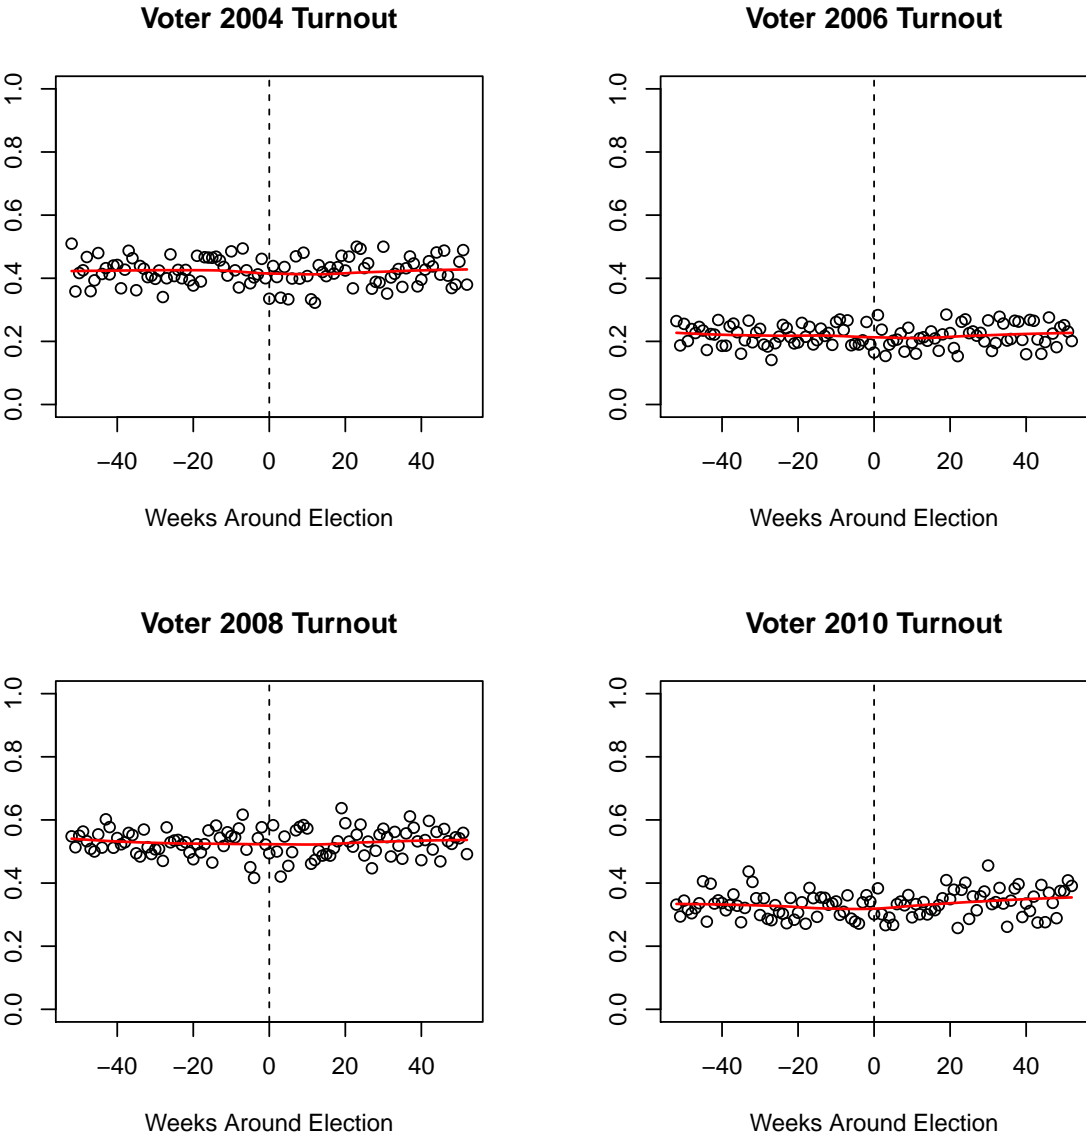


Figure SI1: Prior voting for people seeing household members arrested and charged with misdemeanors during the pre-/post-election periods.

SI4 Simplest Observational Approach

Table SI.1: Basic OLS estimates, including prior vote and voter characteristics

	<i>Dependent variable:</i>		
	Voted 2012		
	(1)	(2)	(3)
HH Member Arrested and Charged	-0.033* (0.004)	-0.009* (0.004)	-0.011 (0.008)
HH Member Convicted	-0.056* (0.005)	-0.016* (0.005)	-0.026* (0.012)
HH Member Sentenced to Jail	-0.100* (0.004)	-0.032* (0.003)	-0.046* (0.007)
2004 Turnout		0.131* (0.001)	0.130* (0.001)
2006 Turnout		0.100* (0.001)	0.100* (0.001)
2008 Turnout		0.368* (0.001)	0.367* (0.001)
Voter Male		-0.017* (0.001)	-0.016* (0.001)
Voter Age (Years)		-0.003* (0.00002)	-0.003* (0.00002)
Constant	0.555* (0.0004)	0.393* (0.001)	0.395* (0.001)
Observations	1,672,616	1,656,185	1,621,932
R ²	0.001	0.223	0.222
Adjusted R ²	0.001	0.223	0.222

Note:

*p<0.05

SI5 Robustness Tests

Logistic Regression This section reproduces Table 2 and Figure 2 from the main paper, estimating the same relationships with logistic regression rather than OLS. Table SI.2 reproduces Table 2, estimating the effect of household member misdemeanor cases, convictions, and jail sentences (that fall within a month of the election) on voting in 2012. As before, we are interested in the interaction between the criminal-justice “treatments” and the indicator for whether they occurred before Election Day. As before, there are substantial effects of household-member conviction and jail sentencing, both in the simplest model (column 1) and one including more background covariates (column 2), with only the estimated effect of conviction statistically significant at $p < .05$.⁶ It is harder to directly interpret these coefficients, but they are similar in magnitude to the OLS estimates. First differences give an intuition: the predicted probability (from Model 1) of voting for someone whose family member was convicted of a misdemeanor after the election is 57%, compared to 43% for someone whose family member was convicted before the election. This difference of 14 percentage points is quite similar to the estimated 17 percentage points reported in Table 2 of the paper.

Figure SI2 reproduces Figure 2, again finding a pattern of short-run demobilizing effects (especially from having seen a household member convicted or jailed before the election) that rapidly taper off into statistical and substantive insignificance. It is worth noting that the y-axis here looks larger than in the original Figure 2, because it presents untransformed logit coefficients, but that the magnitudes of these effects remain similar to those estimated with OLS. In the second panel, for example, the difference in predicted probabilities for voters that saw a conviction in the 40 weeks before the election versus the 40 weeks after (those in the rightmost point on the plot) is about 2 percentage points, quite similar to the (also-non-statistically-significant) estimate in the OLS specification.

⁶As in the main paper, standard errors are clustered at the defendant level.

Table SI.2: Proximal Contact on Voting (Using Case Timing): Logit

	<i>Dependent variable:</i>	
	Voted 2012	
	(1)	(2)
Case Before Election	0.133 (-0.254,0.520)	0.153 (-0.271,0.578)
HH Member Convicted	0.320 (-0.141,0.782)	0.317 (-0.204,0.838)
HH Member Jailed	-0.095 (-0.463,0.273)	-0.063 (-0.496,0.370)
HH Member Convicted * Before Election	-0.690* (-1.337,-0.043)	-0.934* (-1.684,-0.185)
HH Member Jailed * Before Election	-0.509 (-1.038,0.019)	-0.742* (-1.366,-0.117)
Voter Male		-0.046 (-0.267,0.175)
Voter Age (Years)		-0.002 (-0.010,0.006)
HH Member Male		0.150 (-0.140,0.441)
HH Member Black		0.338* (0.051,0.624)
2006 Turnout		0.237 (-0.150,0.624)
2008 Turnout		1.599* (1.309,1.890)
2010 Turnout		1.621* (1.288,1.955)
Constant	-0.029 (-0.317,0.259)	-1.387* (-1.914,-0.860)
	12	
Observations	1,558	1,541
Akaike Inf. Crit.	2,144.568	1,581.240

Note:

*p<0.05

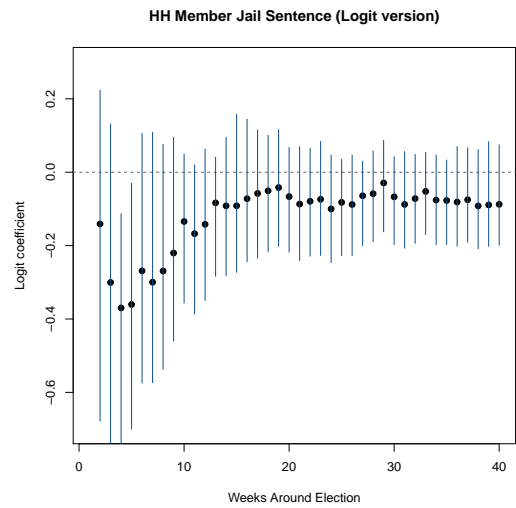
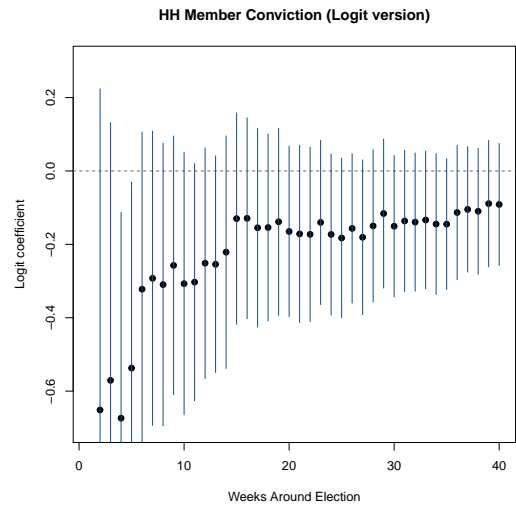
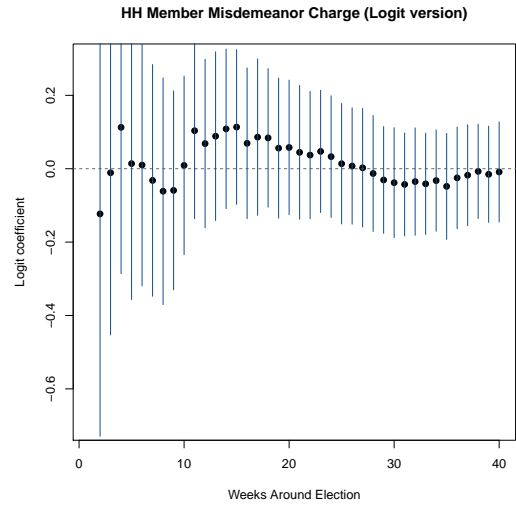
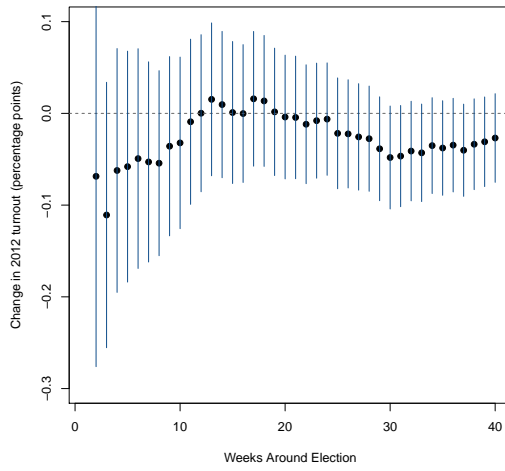


Figure SI2: Reproducing Figure 2 using logistic regression rather than OLS


Match Quality Trying to identify people who experienced proximal contact with the criminal justice system, using administrative data, is not an exact science. We might simply worry that my dataset of “household members” includes many people who have not actually had proximal contact: perhaps defendants’ addresses were recorded wrong, or my geocoding process misplaced them, or residential churn meant that voters and defendants lived at the same address but not at the same time. I discuss match quality in Section SI1, and note there that a hand-coded sample did suggest that most matches were genuine. Nonetheless, even some wrong matches could introduce measurement error. Another way to assuage concerns about mismatches is to restrict the analysis to matches that are even more likely to be correct: cases in which the defendant and the registered voter not only live in the same place, but also share a surname. These represent nearly half of the matches for the time frames used. I show the results from this exercise in Figure SI3: the estimates are substantively similar to those presented in Figure 2, and do not indicate a lasting effect of proximal contact on voting.

Alternatively, we might worry about match quality getting worse over time. After all, I am using a voter file collected as of the end of August 2012. As we get further from the date of the voter file snapshot, matches could get worse as voters and defendants move around. Imagine a match between a voter who appears in the voter file at 10 Singer Lane, Houston (as of 8/26/2012), and a person charged with a misdemeanor in May 2013 who provides the same address: it is possible that the voter no longer lives at that address and is wrongly considered to have proximal contact. If these bad matches get more likely over time, comparing people who had proximal contact before the election (and also closer to the voter-file snapshot date) to those who had it afterwards (further from the snapshot) may generate biased estimates. However, a scenario of this type should bias the estimates in a way that makes them look larger, not smaller (since people without proximal contact generally have a higher baseline probability of voting). I look into this concern further by

Effect of HH Member Misdemeanor Charge on Voting (Varying Windows)



Effect of HH Member Conviction on Voting (Varying Windows)



Effect of HH Member Jail Sentence on Voting (Varying Windows)




Figure SI3: Same analysis as Figure 2, but restricted to households where voter/defendant share a last name.

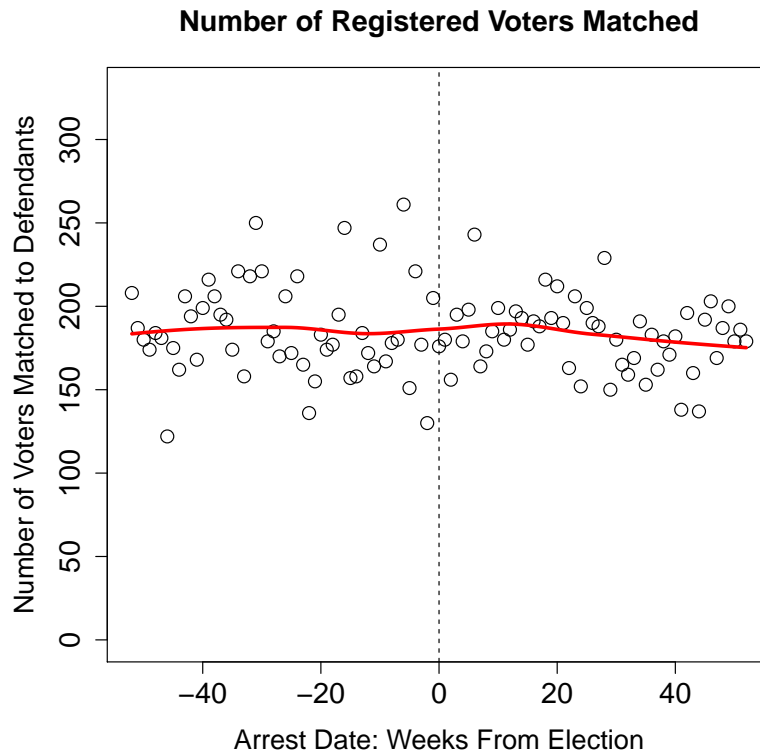


Figure SI4: Binned counts of the number of registered voters matched to misdemeanor defendants each week over time, checking for a decline in match rates after the 2012 election.

plotting, in Figure SI4, the number of voters matched to defendants over time. It doesn't appear that match rates drop off dramatically after the election.

Relatedly, Figure SI5 shows the raw count of people facing first-time misdemeanor cases over time (not restricting to those who matched to registered voters). There do not appear to be any particular time trends in this number either, which is reassuring both on its own (suggesting that we don't see huge seasonality in these cases) and in conjunction with Figure SI4's evidence that household matches aren't diminishing over time.

Statistical Power As noted above, the estimates of proximal contact's effect on turnout are relatively precisely estimated, so it is unlikely that a lack of statistical power is obscuring

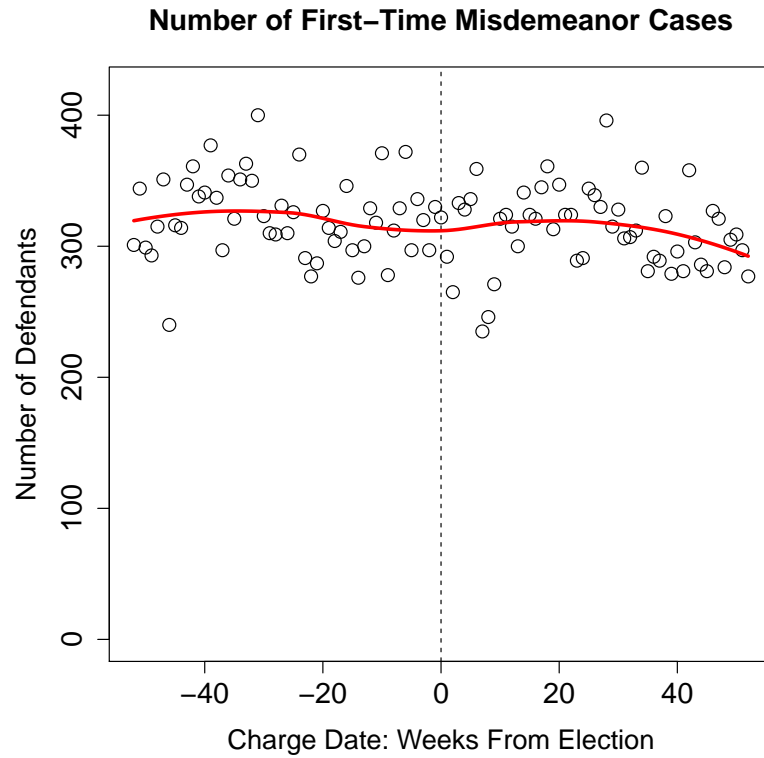


Figure SI5: Binned counts of the number of people charged with their first misdemeanor each week over time (disregarding whether they matched to a registered voter), checking for time trends in filings around the 2012 election.

some large negative effect. It is particularly worth noting that the point estimates presented in Figures 2 and 3 are very near zero, and that the 95% confidence intervals around them do not extend more than a few percentage points from zero.

Sample/Representativeness There are several possible concerns about the way I’ve defined proximal contact for this study. First, it is possible, and even likely, that some of the voters I consider “untreated” have actually experienced some form of proximal contact, through friends or extended family or other non-housemates. Perhaps a voter that I think of as “untreated” because their household member didn’t get arrested until after the election had also seen a friend jailed in September. This is a limitation of measuring proximal contact through administrative data and geographic proximity, rather than self-reports. Still, given existing accounts of how intense the experience of having a loved one incarcerated is, I think we would still expect an additional effect from cases involving household members.

One particular concern, though, is that I might be overlooking people who saw household members charged with felony cases and considering them “untreated”; if they were substantially demobilized by the experience, this could bias my estimates toward zero. This is unlikely to cause major problems, given how few such households there are compared to the population of registered voters, but I check by merging in felony case records and dropping registered voters who experienced felony proximal contact. Figure SI6 repeats the main analysis from the paper with these voters removed; there continues to be no evidence of substantial demobilization outside the few weeks immediately surrounding the election.

However, we might also be curious about these voters who saw their household members charged with felony cases: might they show larger demobilization effects than the main sample, because the economic and social implications of a felony case are larger than those for a misdemeanor? In Figure SI7, I focus on this different sample of voters, replicating the main paper’s analysis. I find no evidence of demobilization; if anything, the point estimates

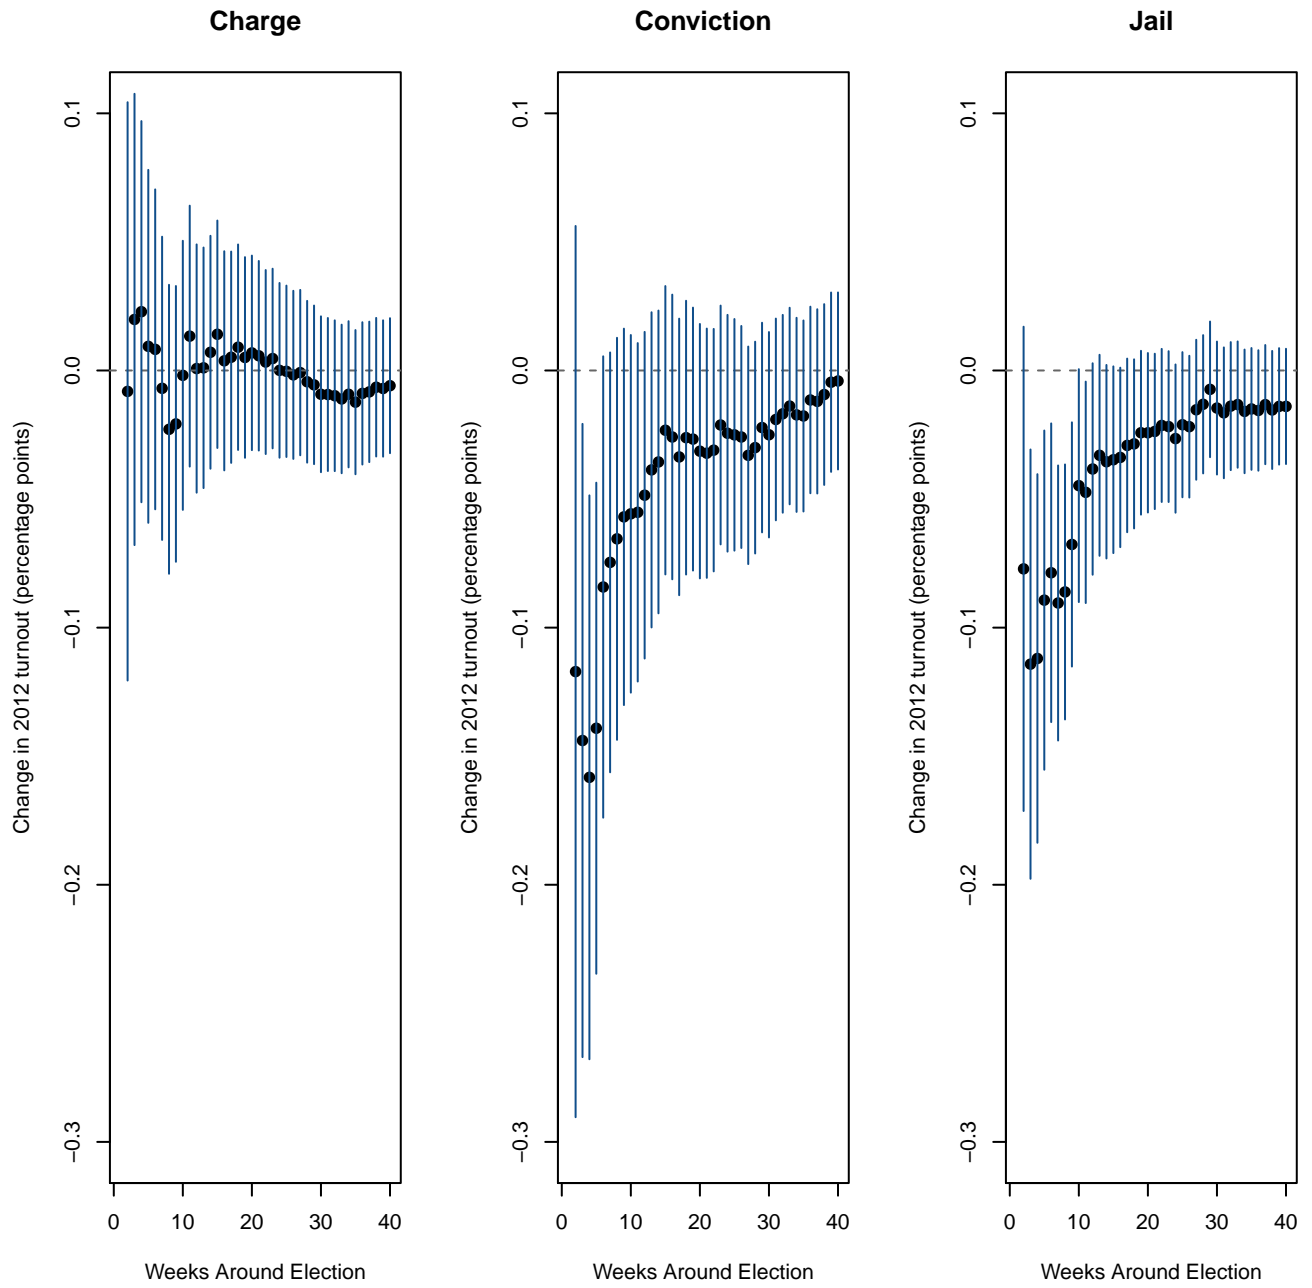


Figure SI6: Effect of proximal contact on household member voting: analogous to Figure 2, but dropping anyone who saw a household member face a felony case.

suggest a small mobilization effect in some cases, though none of these estimates are large or statistically significant. It does not seem that my focus on misdemeanor cases is driving the null effects I find in the main paper.




Figure SI7: Comparable analysis to Figure 2, but estimating the effect of household members' felony cases rather than misdemeanors.

Placebo Test As another check on the validity of the research design, I run a placebo test. Rather than looking at the effect of cases that fall before or after the 2012 election on 2012 voting, I look at the “effect” of cases falling before or after November 5, 2013 (a year after the 2012 election) on 2012 voting. This can serve as a check that something unexpected isn’t happening with the analysis. For example, if there were some sort of seasonality-induced imbalance in cases, it could bias the main results. In this case, given that the main analysis generates null effects, the story could be that there is a true negative effect, but there is positive bias obscuring that effect and resulting in estimates near zero. If that kind of seasonality were occurring, we might also expect to see it in the placebo test. Figure SI8 presents these results; all estimates are substantively small and non-statistically significant, as should be the case for a logically-impossible effect.




Figure SI8: Placebo test: estimating the “effect” of cases falling before/after November 5, 2013 on 2012 turnout.

SI6 Using “Slices” of Time

Rather than continuously broadening the time window used in the analysis, as in Figure 2, I can draw “slices” of the data: that is, rather than moving from one month around the election to two months around the election and so on, I can move from a one-week slice immediately around the election to a pair of weeks that are each one week away from the election date (but on opposite sides). This yields estimates that are based on approximately the same number of observations as I move further from the election, rather than growing more precise as I add more data. Figure SI9 shows this approach with two-week slices of time, estimating the effect of conviction (but not jail time) on 2012 voting.

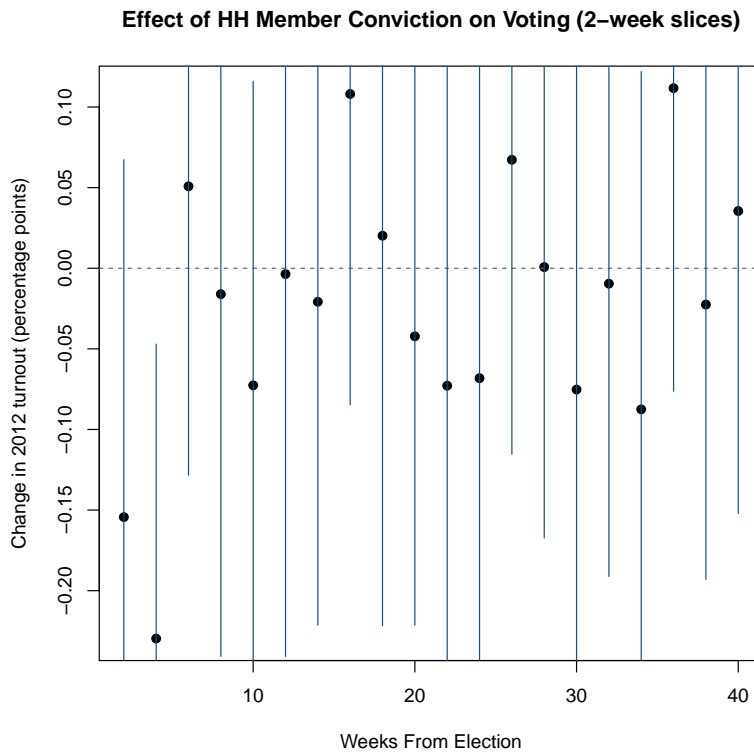


Figure SI9: Effect of conviction on household member voting: analogous to Figure 2, but with a different way of defining time windows.

SI7 Heterogeneity

What if the apparent null effects for most time periods are masking effect heterogeneity? It could be that some household members are demobilized by proximal contact, while others are mobilized, and these effects obscure each other. In this section, I explore the possibility of effect heterogeneity on several dimensions: race, neighborhood, wealth, family relationships, and past voting.⁷ All supporting figures appear below. On the whole, I find little evidence of such heterogeneity.

Race We might expect racial differences in the effects of proximal contact, given immense racial disparities in personal experiences with and views of the criminal legal system (Walker, 2014; Weaver and Lerman, 2014; White, Forthcoming; Maltby, 2017). The voter records used here do not report voters' race, but the criminal court data reports defendants' race, so I use that as a rough proxy for household members' race.⁸ In Figure SI10, I replicate Figure 2 using only the registered voters who were matched to someone who was identified as "Black" in court records. Though these are somewhat more noisily-estimated, the pattern of results remains extremely similar, and does not suggest any lasting effect of proximal contact on voting.

Neighborhood Exposure to Incarceration We might wonder whether some households have already experienced a lot of proximal contact at the neighborhood level, such that they could either show a larger demobilization effect (because this is the last straw), or a smaller one (because this is not providing new information). To explore this possibility, I replicate Figure 2's analyses on households in census tracts with above- and below-median (jail) incarceration rates in the prior year (2011). Again, the results shown (in Figures

⁷This exercise should be viewed as exploratory; I did not envision these analyses when I first conceived of the project, and slicing the data so many ways runs the risk of finding false positives.

⁸These records do not have a category indicating Hispanic or Latino identity, so this analysis is limited to a Black/White comparison.

SI11 and SI12) are similar across the groups, and neither group shows substantial long-run demobilization effects.

Homeownership Earlier in the paper, I discussed both sociopolitical and economic mechanisms by which proximal contact could reduce voting. The economic mechanisms might be more prominent for people who are already economically insecure: the loss of a household member's income or unpaid household work could mean hunger or homelessness. Perhaps the null effects we see are due to cross-cutting effects: middle-class people are somewhat mobilized by the experience, while poorer people are demobilized. I look into that possibility by using one proxy for economic vulnerability: homeownership. This is not a perfect measure of income or wealth, but owning one's home suggests that a temporary loss of household income is unlikely to result in homelessness in the short term. I use tax assessment records from Harris County to identify registered voters who owned their homes as of 2008.

I replicate the main analyses for homeowners in Figure SI15, as well as for the remainder of the sample in Figure SI16(though not being affirmatively matched to a tax assessment record does not perfectly rule out being a homeowner). Again, neither subsample shows lasting demobilization (or mobilization) effects from proximal contact.

Past Voting I also look for heterogeneity over past voting: do people who voted in 2008 show more or less demobilization than those who did not? In Figures SI13 and SI14, I present separate versions of the main analyses for 2008 voters and non-voters. Again, the pattern of null effects in the medium to long term is consistent across both groups.


Family Structure Finally, I look at people with different relationships to the household member charged: some may be parents watching a child interact with the system, while others are watching a romantic partner, and others a roommate or other household member. I try to identify parents and romantic partners of defendants using the age and gender of

household members. In this crude (and heteronormative) approach, I assume that people within fifteen years of age of one another and of opposite sexes are romantic partners, while people more than fifteen years older than the defendant are a parent or other older relative. When I run the main analyses for voters assumed to be parents of the defendant, the estimates again indicate no long-term effect of proximal contact (shown in Figure SI17).

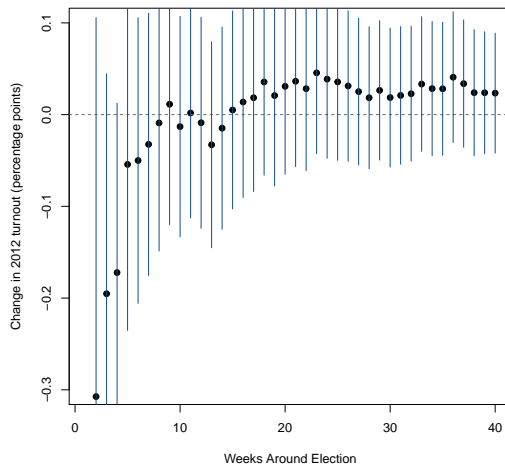
The estimates for the assumed romantic partners of defendants (shown in Figure SI18) are somewhat more noisily-estimated and varied: there is no long-run effect of seeing one's household member face a misdemeanor charge or conviction, but the estimates for certain time windows suggest a moderate negative effect of jail on voting, even months away from the election. As noted above, these results could be spurious; slicing the data in many ways generates a substantial multiple-testing problem. Still, they deserve further testing; future work could examine whether this result holds among voters who are married to the arrested household member. Persistent demobilization effects among people who are perhaps the most affected (both economically and socially/emotionally) by the arrest of a household member could suggest an important distinction within the broader experience of "proximal contact."

Black Voters

Effect of HH Member Misdemeanor Charge on Voting (Varying Windows)



Effect of HH Member Conviction on Voting (Varying Windows)



Effect of HH Member Jail Sentence on Voting (Varying Windows)




Figure SI10: Same analysis as Figure 2, but restricted to registered voters matched to someone whose race was recorded as “Black” in court records.

High- and Low-Incarceration Neighborhoods

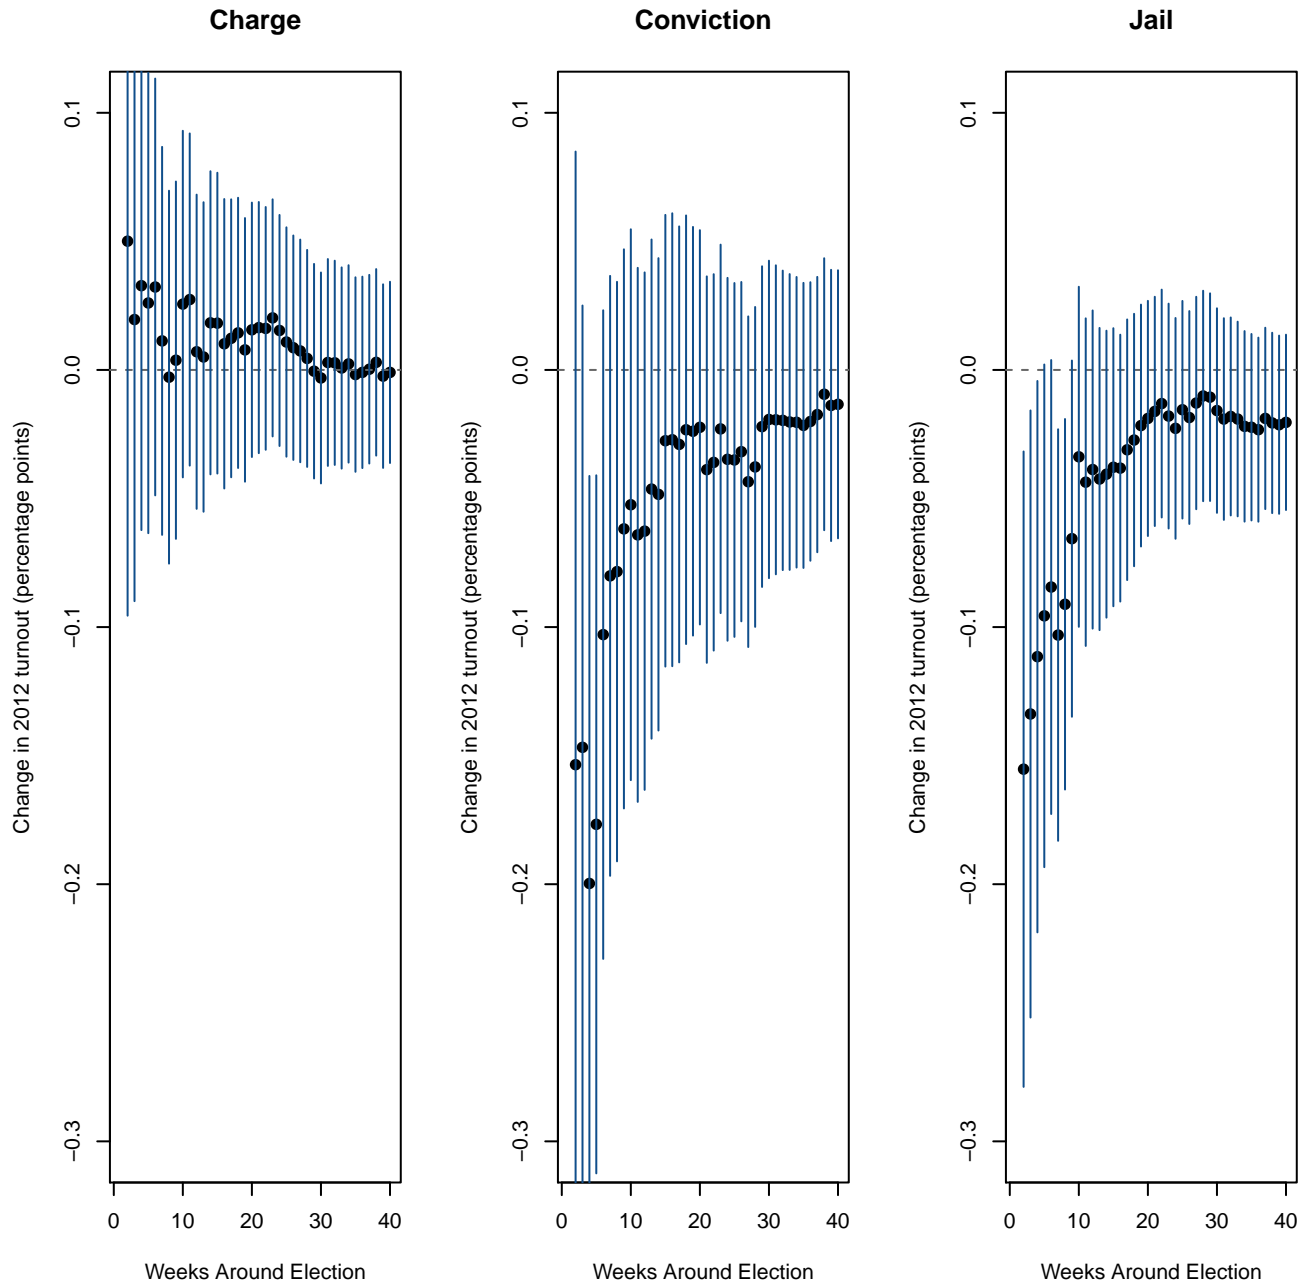


Figure SI11: Same analysis as Figure 2, but restricted to registered voters living in neighborhoods with below-median incarceration rates.

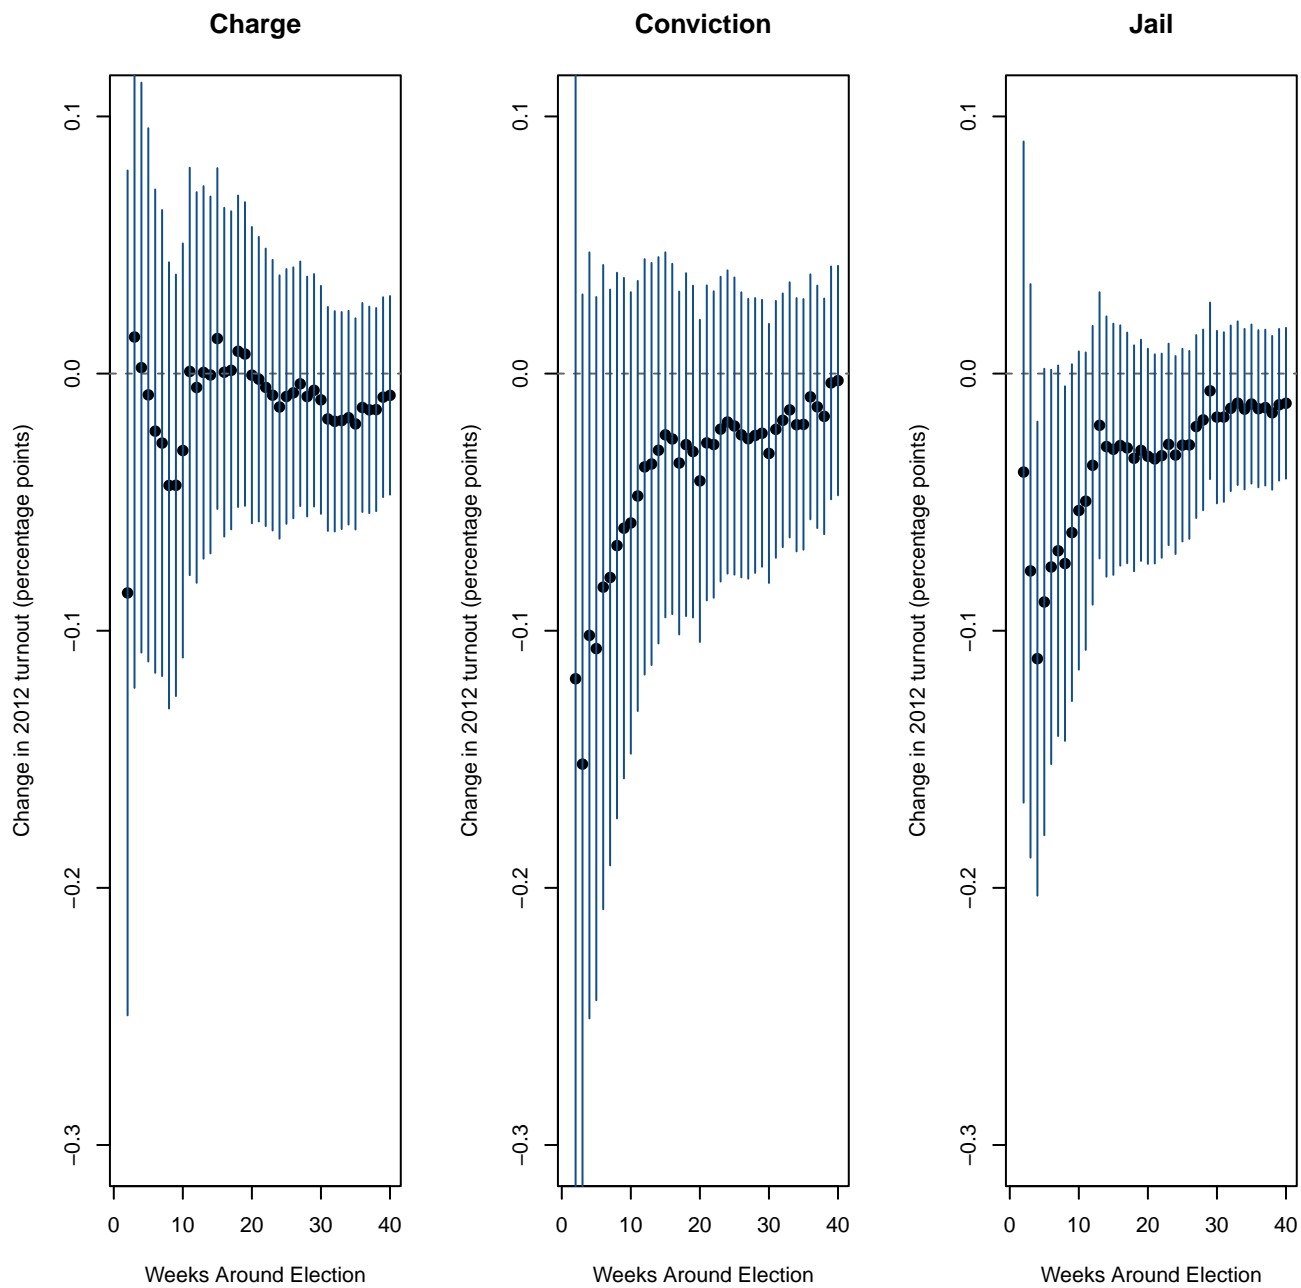


Figure SI12: Same analysis as Figure 2, but restricted to registered voters living in neighborhoods with above-median incarceration rates.

Past Voting (2008)

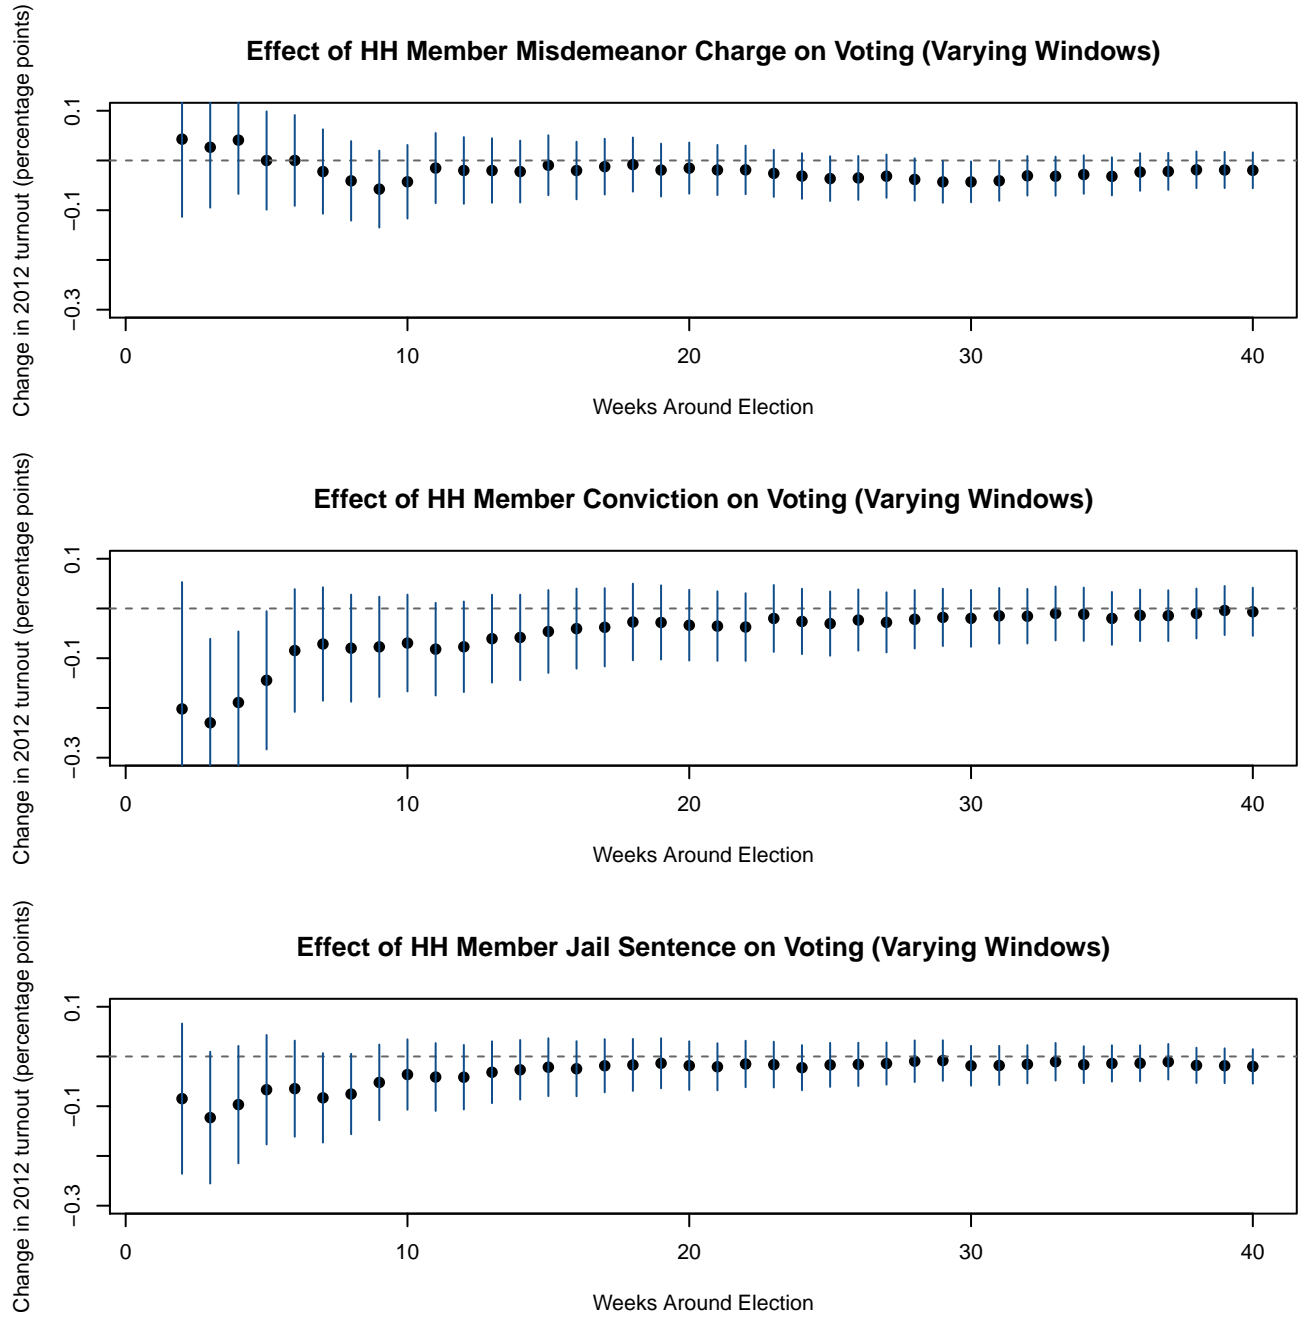


Figure SI13: Same analysis as Figure 2, but restricted to registered voters who are recorded as having voted in 2008.

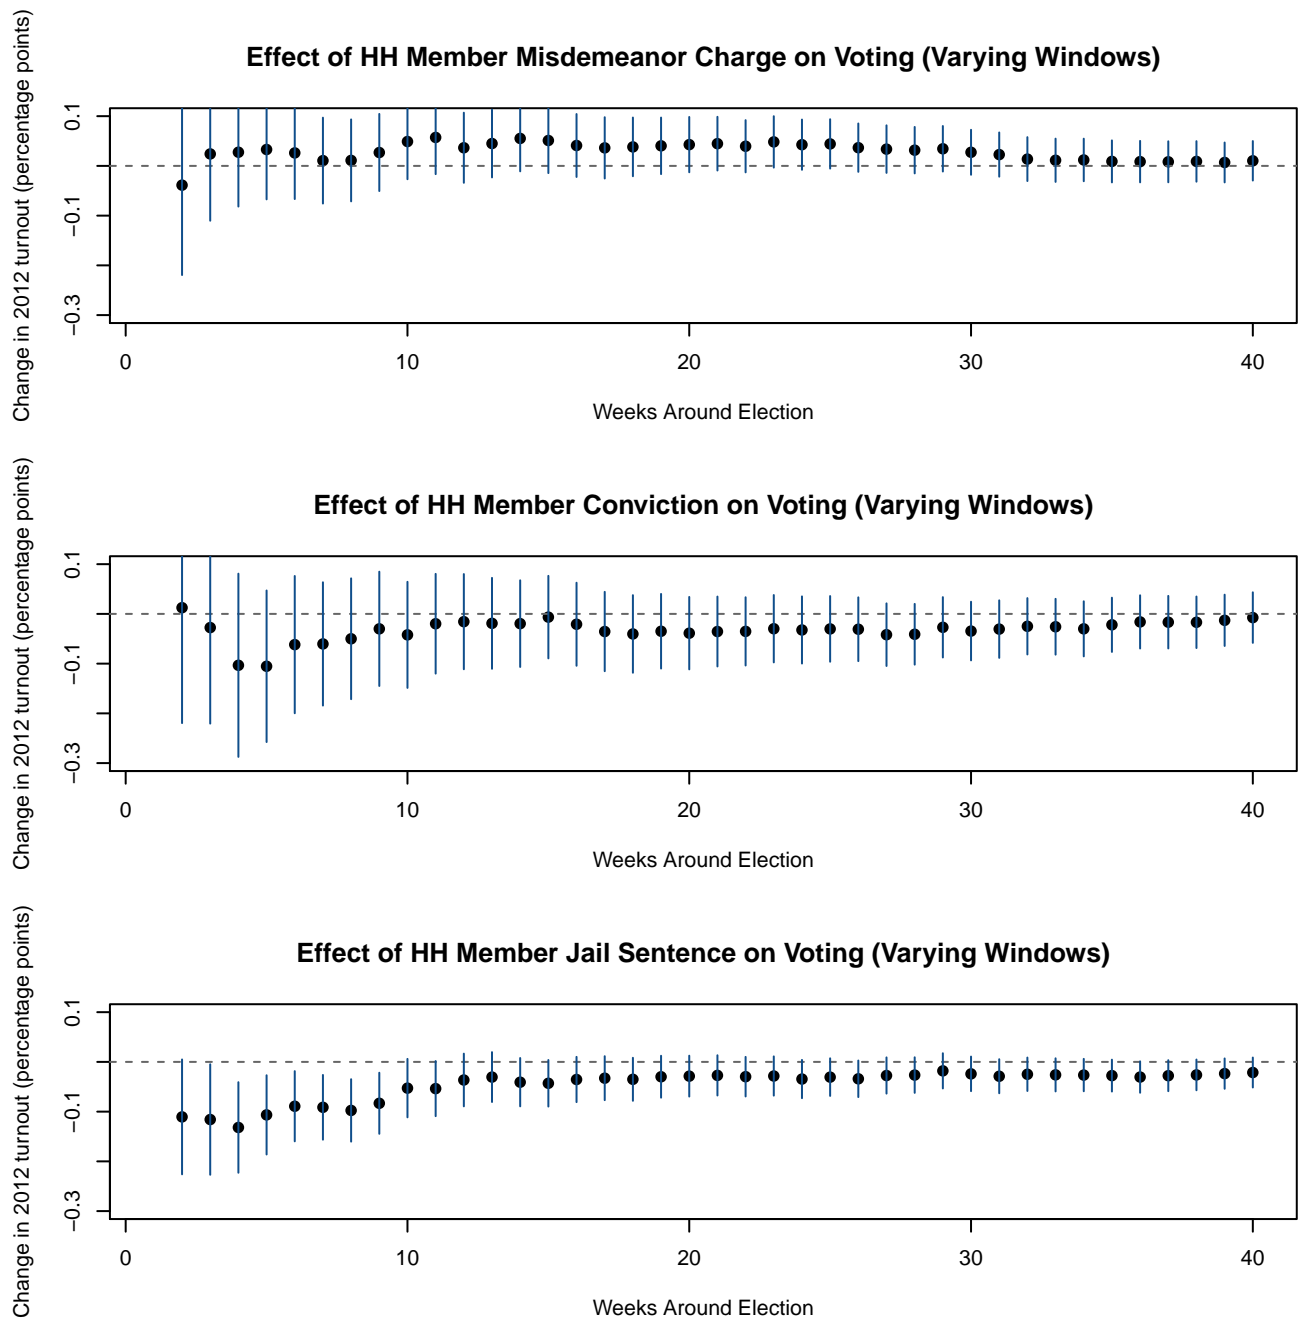


Figure SI14: Same analysis as Figure 2, but restricted to registered voters who are not recorded as having voted in 2008.

Homeowners

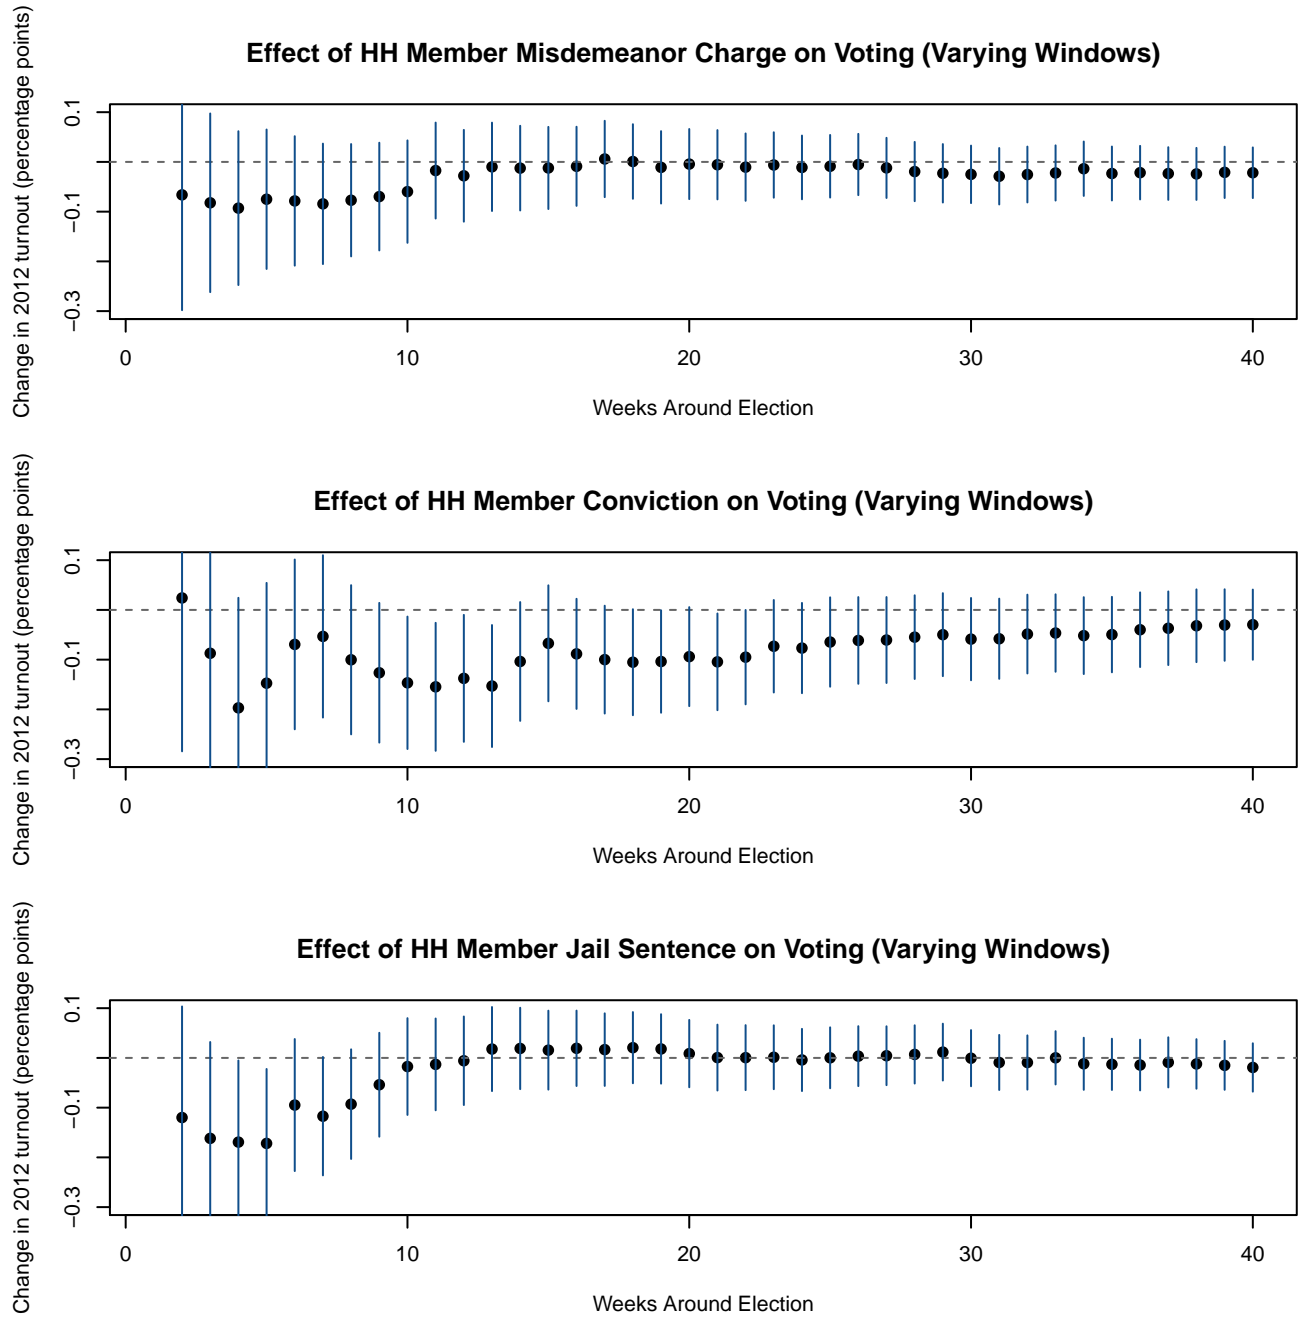


Figure SI15: Same analysis as Figure 2, but restricted to registered voters who appear to own their homes.

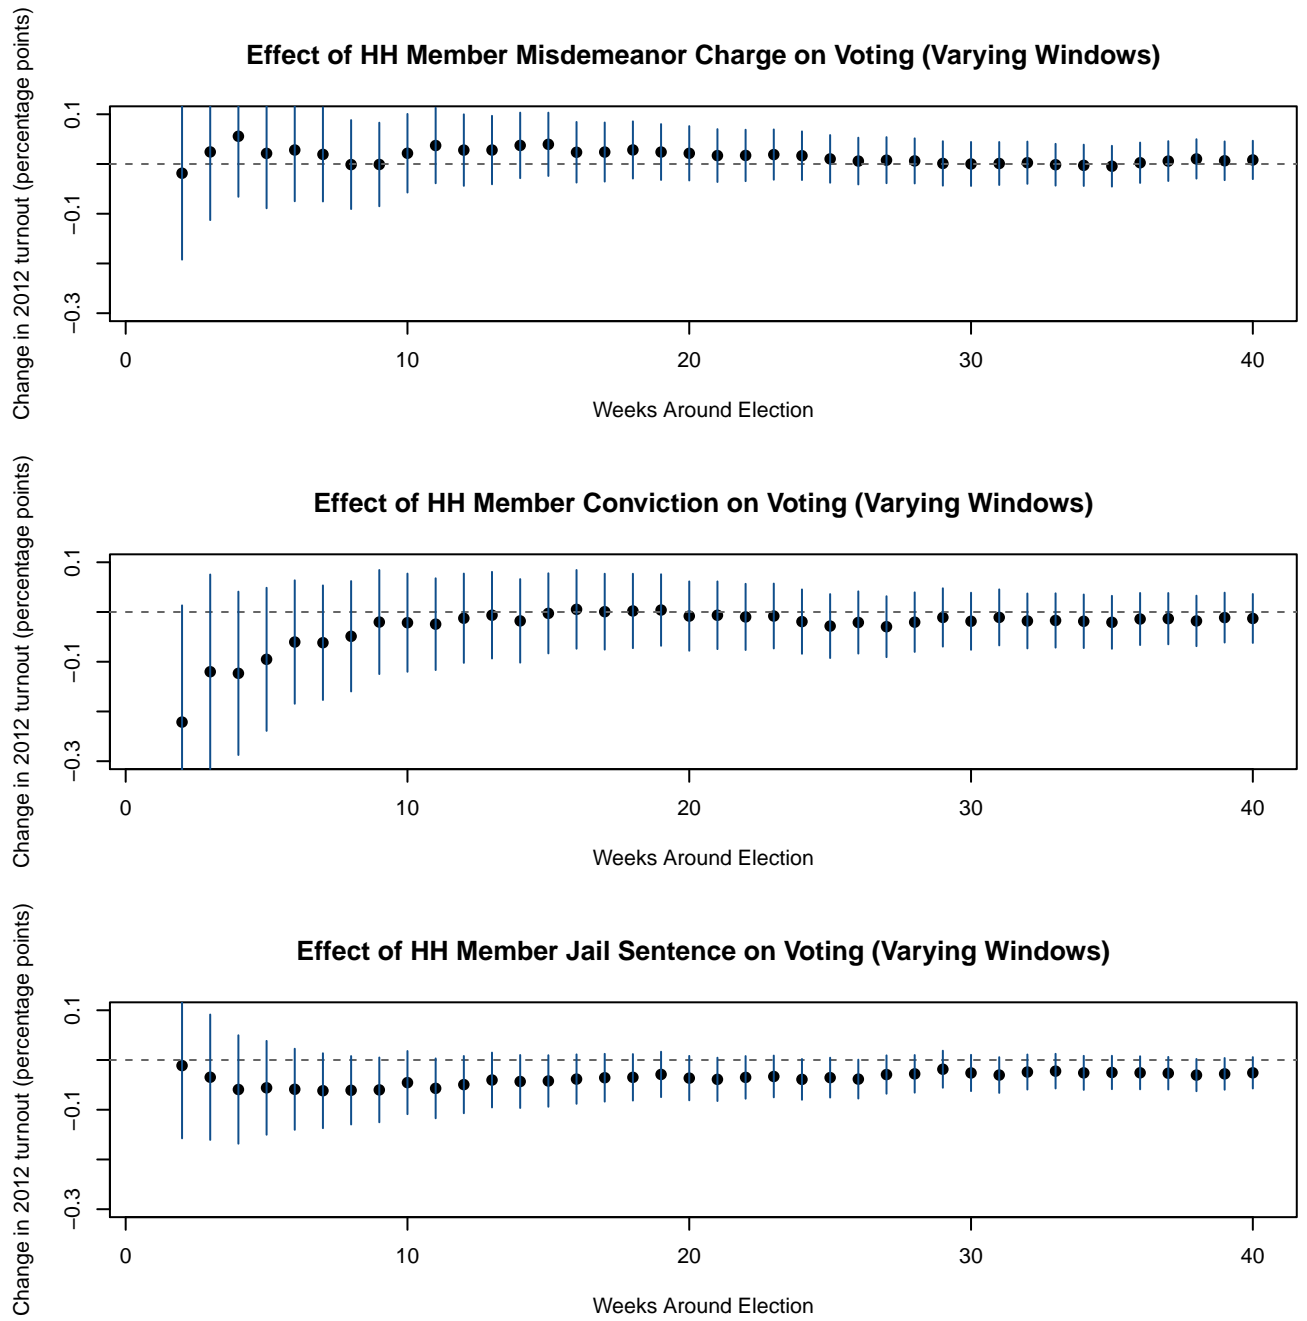


Figure SI16: Same analysis as Figure 2, but restricted to registered voters who don't appear to own their homes.

Parents and Partners

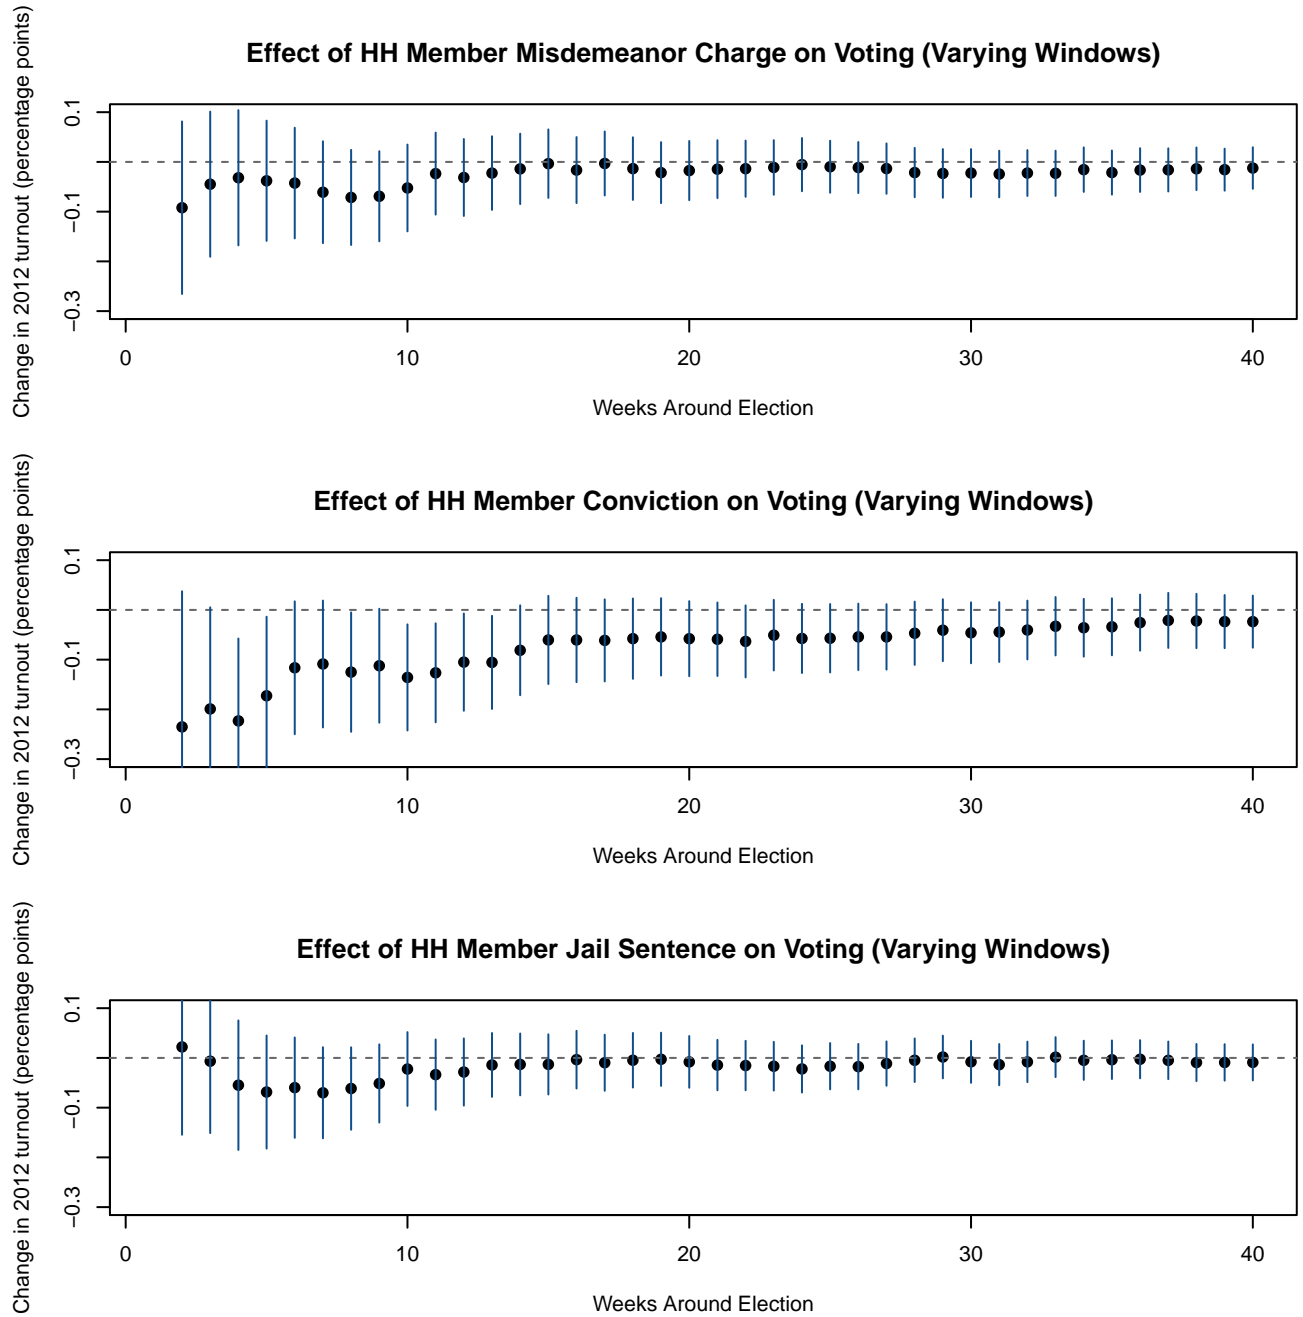


Figure SI17: Same analysis as Figure 2, but restricted to apparent parents/older relatives of the person charged.

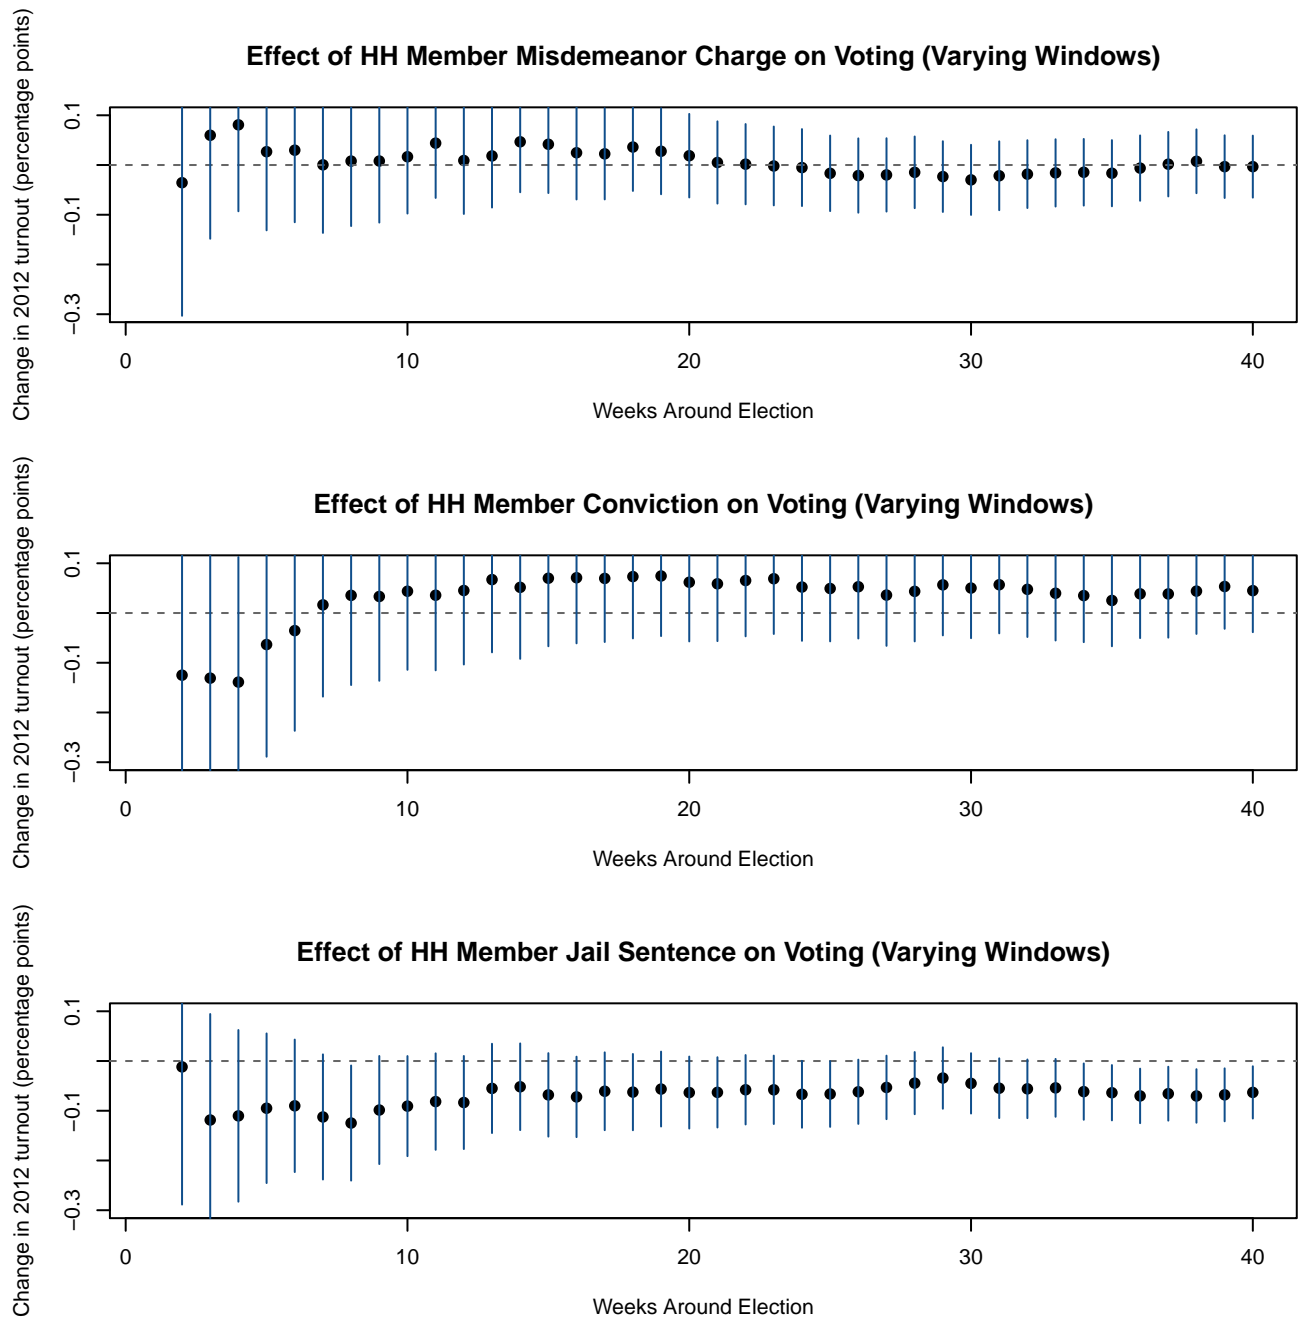


Figure SI18: Same analysis as Figure 2, but restricted to apparent romantic partners (similar age, opposite sex) of the person charged.

SI8 Another Look at Short-Term Effects

This section presents a set of plots that display the short-term demobilizing effects of household member jail time in a different way. Rather than estimating a difference in means between voters that see cases before and after the election at varying bandwidths, as the main paper's Figure 2 does, here I simply present voter turnout within weekly bins. Figure SI19 plots the proportion of people that voted, based on what week they saw their household member charged. The red lines are Loess smoothers. Consistent with the main Figure 2, this plot suggests that people experiencing cases in the few weeks just before an election become less likely to vote, but the effect does not persist among people who saw cases in the months prior to election day.

We might also wonder whether people who don't vote in 2012 due to a household member's criminal case become less likely to vote in 2016: does the pre-election drop in voting shown in Figure SI19 persist into the next election? Figure SI20 explores this question by merging a 2017 copy of the voter file from Harris County into the main dataset and remaking the plot with 2016 turnout. This approach merits some caution, as people may have moved or dropped off the voter file for a range of reasons. But the plot doesn't suggest a lasting demobilization, even among those people who saw cases right before the election and became less likely to vote in 2012. The red curves are much flatter than in Figure SI19.

2012 Turnout by Household Members

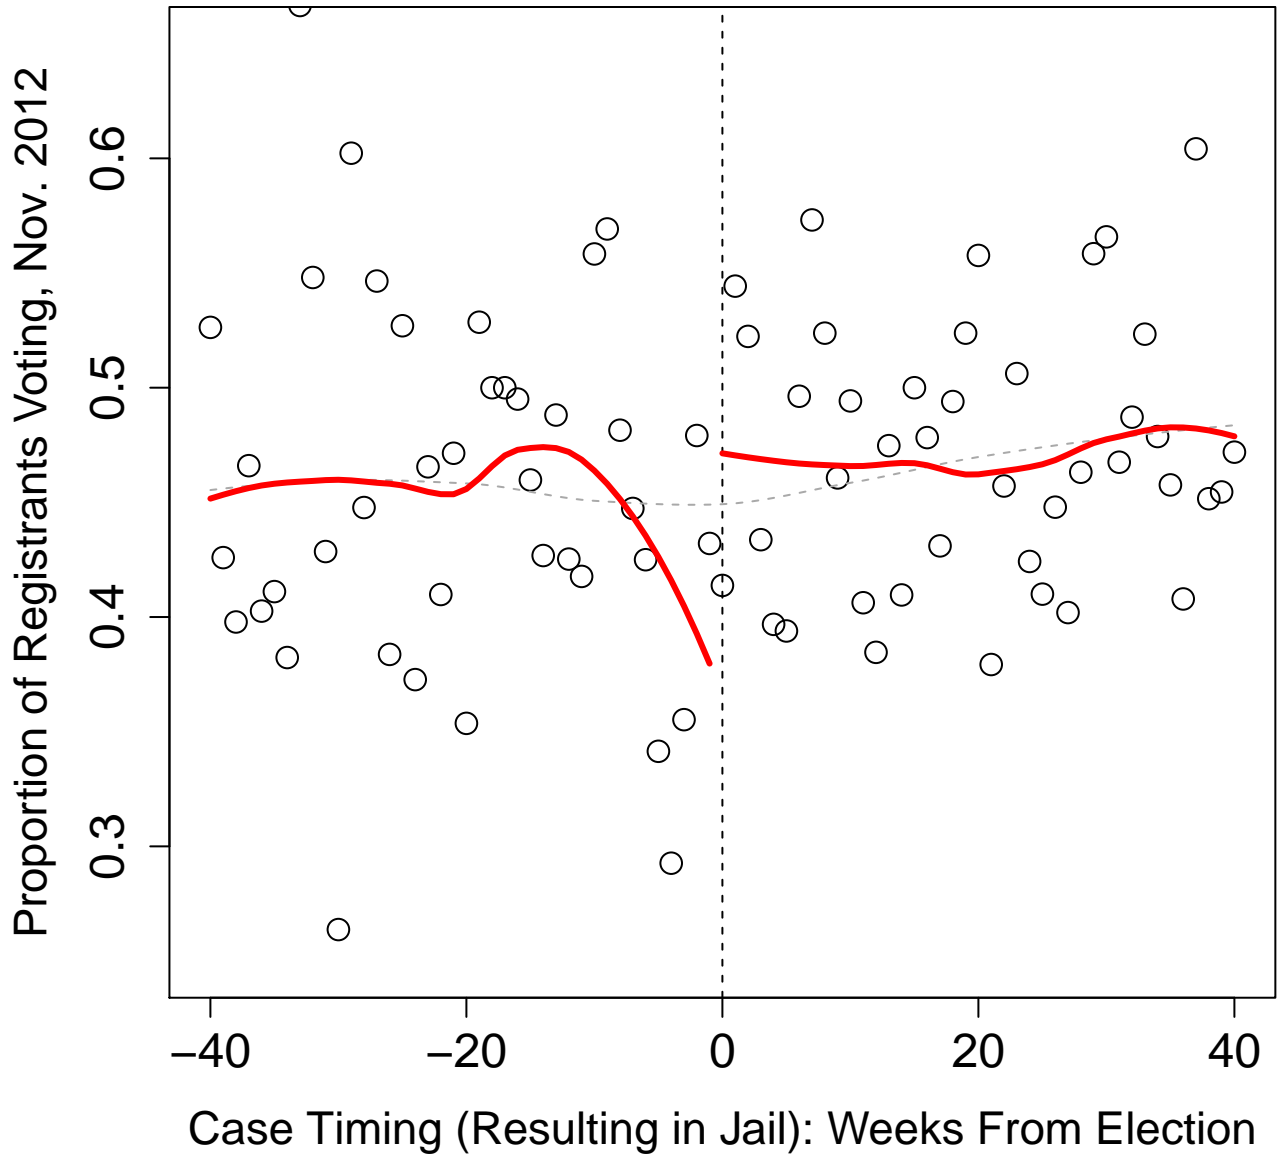


Figure SI19: Voter turnout by week of household members charge

2012 Turnout by Household Members

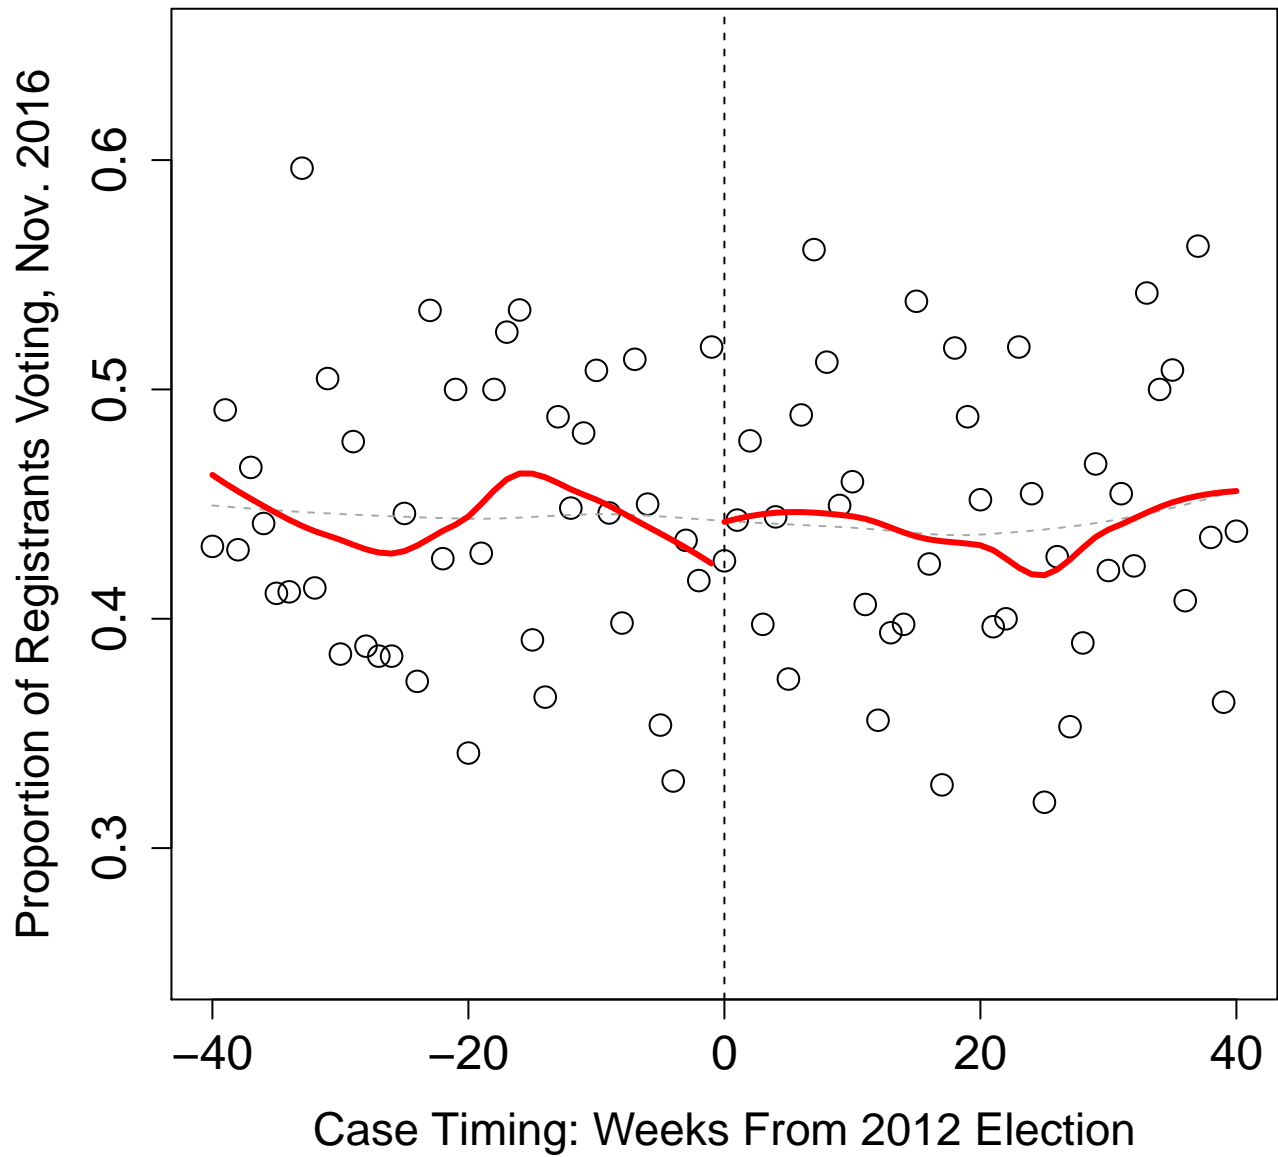


Figure SI20: Analogous to Figure SI9, but looking at 2016 voting to see whether short-run demobilization in 2012 persists to the next presidential election.