

# IDEOLOGICAL DIVERSITY IN SOCIAL INTERACTIONS AND POLITICAL POLARIZATION

Kirsten Cornelson\*

October 16, 2020

## Abstract

Can socializing with people who disagree with you reduce political polarization? I answer this question using a shock that induces us to socialize and discuss politics with a more ideologically diverse set of people: Thanksgiving. I use both Canada/U.S. comparisons (exploiting the fact that Thanksgiving occurs at different times in the two countries), as well as a series of Canadian electoral reforms affecting the timing of elections, to show that people converge towards the political viewpoints of their families after Thanksgiving and that this reduces polarization. These results suggest that increasing political diversity within our social circles can reduce polarization. JEL codes: P16, Z1

---

\*Department of Economics, University of Notre Dame, 3051 Jenkins Nanovic Hall, Notre Dame, IN 46556 e-mail: [kcornels@nd.edu](mailto:kcornels@nd.edu). I thank Taryn Dinkelman, Kasey Buckles, Ethan Lieber, Dan Hungerman, Abigail Wozniak, and participants at the Mid-Midwest Applied Microeconomics Workshop for their helpful comments.

# 1 Introduction

Political polarization is increasing in both the U.S. and Canada. In the U.S., election surveys show that the policy preferences of Republicans and Democrats have been diverging since the 1970's (Abramowitz and Saunders, 2008). The level of hostility towards members of the other party has also risen substantially over the same period (Gentzkow, 2016). The partisan divide in Canada has been increasing at a similar rate since the early 1990's (the earliest time period for which data is available), both in terms of policy preferences (Kevins and Soroka, 2018) and hostility (Boxell et al., 2020).

Popular and academic accounts of increasing polarization have focused heavily on two potential explanations: the introduction of partisan cable news channels, and the rise of social media (e.g., Sunstein, 2017; Taibbi, 2019). Careful research has shown that these explanations have some merit. For example, researchers have used the gradual roll-out of Fox News as well as channel position instruments to show that the network does indeed pull its viewers to the right (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017). Experiments on Facebook have shown that deactivating Facebook, or getting users to subscribe to a greater variety of media sources, reduces polarization (Hunt Allcott and Gentzkow, 2020; Levy, 2020). However, while the media and the internet clearly play some role in increasing polarization, they cannot fully account for the increase in polarization over time, which predates either phenomenon.

While increases in ideological segregation online and in the media have received a great deal of attention, a similar and parallel phenomenon has gone relatively unnoticed: an increase in the intensity of our real life “echo chambers”. As noted by Gentzkow and Shapiro (2011), our real life social circles are substantially more ideologically segregated than our media or internet consumption. While it is difficult to directly track how this level of ideological segregation has changed over time, we can infer something from changes in *who* we are spending our time with. Table I shows the number of minutes per day that Americans spent with different groups in 1965, and over the period 2010-2018.<sup>1</sup> The table shows that we increasingly choose to spend time with the people we are ideologically closest to. People who are married spend a full 2 hours more per day with their spouses than they did in the 1960's, and have also increased the amount of time they spend with children and other household members. As we move to more distant groups, the direction of this change reverses. We spend exactly the same amount of time with non-household relatives and friends as we did in 1965 (100 minutes per day, on average), but less time

---

<sup>1</sup>Unfortunately, no data exist that allow me to make the same comparisons for Canadians.

socializing with coworkers outside of work (13 minutes per day, compared to 20 minutes in 1965), with neighbors (7 minutes per day, compared to 28 minutes in 1965), or in the presence of any other non-household person (around 36 minutes per day, compared to 65 in 1965.) Because our spouses and closest friends are the most similar to us politically<sup>2</sup>, the net effect of these time changes is almost certainly to reduce the diversity of political opinions we encounter in our day-to-day lives.

Can increasing the ideological diversity of our social interactions help to reduce polarization? If so, how large is this effect? In this paper, I answer these questions by using a shock that induces us to socialize and discuss politics with a wider variety of people: Thanksgiving. I use two empirical strategies to measure the effect of Thanksgiving on both exposure to conflicting opinions and political polarization. First, I take advantage of the fact that Thanksgiving occurs at different times in the U.S. and Canada. I ran a repeated weekly survey of MTurk workers in the autumn of 2019 in which I asked questions about political conversations and political beliefs, using Americans as a control group around Canadian Thanksgiving and Canadians as a control group around American Thanksgiving. This survey shows that Thanksgiving does represent a shock to ideological information. In non-Thanksgiving weeks, respondents are most likely to discuss politics with their spouses and friends, who are also the groups that they report are closest to them in terms of political beliefs. During Thanksgiving, there is a large increase in the probability of discussing politics with siblings and extended family members, who are more politically dissimilar to the respondents.

Next, I show that the additional interaction with our extended family induced by Thanksgiving affects our political beliefs. Specifically, I show that respondents who have siblings/extended families to the right of them on a political orientation scale move significantly rightward on an opinions index immediately after Thanksgiving, while respondents whose families are to the left of them move left (see Figure III and Figure IV.) The effect is strongest in the week after Thanksgiving, but is discernible for at least 3 weeks after that. As a placebo test, I show that the effect does not occur among French Canadians, who do not celebrate Thanksgiving. This is particularly important in ruling out any effect of the Canadian election in October of 2019. The effect is also not driven by differential selection into answering the survey during the week of Thanksgiving.

A potential concern when interpreting these results is that Thanksgiving might operate through some channel other than through interaction with family. For example, perhaps

---

<sup>2</sup>I show this below

people become happier or more thankful around this time and move to the center as a result; this would produce similar patterns to the “social interaction” hypothesis. To support the interpretation that the effects of Thanksgiving work through changes in social interactions, I present the results of two additional exercises. First, I show that people with identical political beliefs before Thanksgiving move in different directions depending on whether they were exposed to a more left- or right-wing family; this is inconsistent with a general depolarization. Secondly, there is no similar “effect” for family members or friends for whom interaction is unaffected by Thanksgiving. Both of these patterns suggest that the effect of Thanksgiving on political beliefs operates through its impact on the tendency to socialize with siblings and extended family.

How large is the effect of Thanksgiving on political opinions? Using my regression results to predict the change in political opinions for members of my sample suggests that Thanksgiving reduces political polarization by around 0.12 standard deviations in the week following Thanksgiving. To put this number in perspective, Hunt Allcott and Gentzkow (2020) show that deactivating Facebook for 4 weeks reduces a similar measure of polarization by 0.16 standard deviations.

While the survey shows that Thanksgiving has a noticeable effect on political opinions, it leaves two important questions unresolved. First, how general is this effect? Perhaps MTurk workers, who are disproportionately young, educated, and left-wing, are more susceptible to being influenced than the general population. Secondly, does this change in political beliefs correspond to any change in political behavior? It is possible that respondents are simply more likely to *report* different political beliefs after Thanksgiving, either because they view these beliefs as more socially acceptable, or because different viewpoints have become more salient.

To address these questions, I next use a series of Canadian political reforms that affected the timing of elections to examine whether there is a change in voting behavior when an election takes place right after Thanksgiving. During the mid-2000’s, nearly all Canadian provinces and the federal government moved from floating-date elections, held at the behest of the ruling party, to fixed-date elections, held on a specific date every 4 years. Three jurisdictions selected fixed dates that occurred within 1 month after Thanksgiving, while the rest chose dates either just before Thanksgiving, or in the spring. I use a difference-in-difference specification comparing voting behavior in the post-Thanksgiving group to other jurisdictions, before and after the reforms.

If Thanksgiving operates in the way I have suggested, then more voters should occupy

a jurisdiction's political "center" - the modal political viewpoint - immediately following Thanksgiving. This should translate into more votes for the winning party, and less dispersion in votes overall. To capture this logic, I examine the impact of Thanksgiving on measures of inequality in vote shares. Perfect equality of vote shares across parties in an election represents the maximum degree of disagreement within the electorate, while perfect inequality (all votes going to one party) represents the maximum degree of agreement. I therefore examine whether holding elections after Thanksgiving results in increased inequality in vote shares, or agreement among the electorate.

Figure VI shows the results of this exercise. Focusing on the top left panel for now, the graph shows that in the jurisdictions that adopted fixed dates after Thanksgiving, the winning party's vote share increased by nearly 13 percentage points after the reform compared to the control group. In other words, there was substantially more agreement among the electorate once elections were moved to the month after Thanksgiving. As shown in the other three panels, similar patterns hold for other measures of inequality in vote shares.<sup>3</sup> Again, this is a large effect compared to what has been documented in the previous literature. For example, Martin and Yurukoglu (2017) show that watching Fox News for one hour per week for a year increases the Republican vote share by around 7 percentage points.

In sum, my results show that socializing with people who disagree with you can reduce polarization, and that the effects are relatively large. One dinner with your extended family has a similar effect to deactivating Facebook for a month, or watching Fox News for one hour per week for a year. More broadly, the analysis also suggests that changes in our social habits may have played an important role in increasing polarization over time, and that reversing these changes may help combat polarization.

Of course, it is important to emphasize that my results capture the effect of *marginally* increasing ideological diversity in social interactions. In particular, they do not imply that placing two randomly selected people in a room together is likely to be an effective way to reduce political disagreement. Instead, they show that we are willing to change our minds when we spend more time with people we like, trust, and have something in common with - but disagree with nonetheless. Fortunately, this describes most of the people that we might encounter by slightly broadening our social circles. Secondly, my results capture the effect of spending time with people we don't see often in real life. If the first hour with someone has a bigger impact than the tenth or the hundredth, my results

---

<sup>3</sup>I explain these measures in more detail in a later section.

will be larger than the average treatment effect of inducing people to spend more time with their extended families, neighbors, or coworkers on a regular basis. Nonetheless, the large magnitude of this “first hour” effect suggests that inducing people to socialize with more political diverse people even slightly more often may be very effective in reducing polarization.

## 2 Data and empirical strategy

### 2.1 Canada-U.S. comparison

To ascertain the effect of Thanksgiving on exposure to different viewpoints and on political opinions, I begin by comparing a set of American and Canadian survey respondents around American and Canadian Thanksgiving. I examine whether individuals have political discussions with a more ideologically diverse group of people around Thanksgiving, whether these discussions are more heated, and whether individuals change their political views after Thanksgiving. I do this by comparing responses in the weeks during or just after Thanksgiving to responses from people within the same country in different weeks, and people from the other country (who are not experiencing Thanksgiving) in the same week. I provide more details on the empirical strategy below, after describing the data collection process and outcome variables.

#### 2.1.1 Survey sample

I recruited a set of approximately 1,000 MTurk respondents in the fall of 2019. The first set of respondents were recruited during the week of October 5th (9 days before Canadian Thanksgiving), and answered surveys weekly for the following 6 weeks. The second set of respondents was recruited during the week of November 17th (11 days before American Thanksgiving), and also answered surveys weekly for the following 6 weeks. Because I will be using ratings of the political beliefs of self and family throughout the analysis, I drop respondents who do not have information on these variables.<sup>4</sup> Most respondents did not answer surveys in every week; the average number of weeks in the sample is 4.1. The final number of respondent-week observations is 3,859.

---

<sup>4</sup>Of an initial 1,023 respondents, 16 did not have information on their own left/right orientation. Of the remaining 1,007 respondents, a further 74 did not report left/right orientation for either siblings or extended family, who are the groups I use to construct the family political orientation variable.

### 2.1.2 Baseline information

In each respondent’s first survey wave, I collected baseline information on demographics, as well as information about the political orientation of the respondent. To measure political orientation, I asked respondents to place themselves on a left-right scale, ranging from “extremely left wing” (coded as 0) to “extremely right wing” (coded as 6). Table II shows summary statistics on both demographics and political orientation. The first row shows that the sample is approximately 1/3 Canadian. This is because I was unsuccessful in recruiting the desired number of participants from Canada.<sup>5</sup> The next rows show that the MTurk sample is disproportionately young, with approximately 2/3 of workers in both countries falling in the 26-45 age range. The sample is also approximately 58% male. The racial composition of respondents is close to the population distribution within each country (although slightly less likely to be black and more likely to be Asian within the U.S. sample), and is far more educated than the population average. The sample is also left-wing, with an average rating of 2.5 on a scale of 0 to 6. The non-representativeness of the MTurk sample is one of the key reasons why I also present results using Canadian election data below.

Next, I gathered information on the political viewpoints of different groups of people that the respondent might speak to about politics. I use this information to measure whether respondents are exposed to a different set of viewpoints during Thanksgiving, as well as to measure the expected direction of change in political opinions after Thanksgiving. I asked respondents to rate the views of the following people/groups, using the same scale as in their self-ratings: spouse or partner, parents, siblings, adult children, extended family, friends, coworkers, and the people in the respondent’s neighborhood. Respondents were instructed to think of a typical or average member of each group when the question referred to more than one person. The second column of Table III shows the average level of disagreement between the respondent and each group, where disagreement is measured as the absolute difference between the respondent’s own 0-6 rating and the rating they gave for each group. Respondents are closest politically to their spouses and friends, with an average difference of about 0.8 points and 1.0 points respectively. They report being furthest from coworkers/neighbors and extended family, with an average difference of around 1.6-1.7 points. Siblings, parents, and adult children fall somewhere in between.

As I will show below, Thanksgiving primarily increases the probability of discussion

---

<sup>5</sup>MTurk workers outside of the United States are paid in Amazon gift cards rather than in cash, which may account for the smaller worker pool in Canada.

with siblings and extended family. To be more concrete about the nature of this shock, Figure II shows the distribution of political opinions on the 0 to 6 scale for the MTurk respondents, and for their families (the average of siblings and extended family.) The MTurk sample is shifted to the left, with modal respondent rating their political views as “left-wing” (a 1 on a scale of 0 to 6.) In contrast, the reported views of their families are normally distributed with a mean of 3.1, which corresponds to “neither left nor right wing” on the scale. Within the survey, therefore, we should think of the shock as taking a group of left-wing respondents and exposing them to a more ideologically representative segment of the population.

### 2.1.3 Weekly survey questions

During each week of the survey, I asked respondents whether they had discussed politics or current events in the previous week, and if so, with which people. The first column of Table III shows the probability that a respondent reported having a discussion with each group of people in the previous week, averaged across all survey weeks.<sup>6</sup> 86% of respondents report having a political discussion with anyone in the past week. The most common conversation partners are friends (46.7%) and spouses/partners (45.0%), followed by parents (33.5%) and coworkers/other (25.1%). Siblings, extended family, and adult children have the lowest rates, at 7.9-19.4%.

In order to examine whether discussions become more disagreeable around Thanksgiving, I also ask respondents questions about their most recent political discussion within the past week. I ask them to rate the level of agreement among participants on this discussion on a scale of 1 to 10, where “1” indicates that they disagreed on almost everything while “10” indicates that they agreed on almost everything. On the whole, discussions appear to be quite agreeable: the mean value of this variable is 7.6. I also ask respondents how much they enjoyed this discussion, where the possible responses are “not at all”, “somewhat”, or “a lot”.

Finally, I take a weekly measure of respondents’ opinions on political issues. Unlike the political *orientation* rating, which each respondent reports only once at baseline, the opinion index captures the respondent’s left/right position in a way that varies over time. In the first week and every week thereafter, I show each respondent a set of 8 statements, randomly chosen each week from a pool of 20, and ask them to rate their agreement on

---

<sup>6</sup>These probabilities are not conditional on having the relationship described; for example, uncoupled people are included as zeros in the variable “discussed with spouse”.



a scale of 1 to 5 (from “strongly disagree” to “strongly agree”). The statements were chosen to reflect typical left- and right-wing positions on a variety of issues, as classified by political scientists at the Comparative Manifesto Project (Krause et al., 2020). An example of a left-wing statement is: “The military budget would be better spent on social programs like health or education.” An example of a right-wing statement is: “Spending on social programs in this country has gotten out of control. Individuals should not expect the government to provide for all of their needs.” The full set of statements and more details about the development of this portion of the survey are available in the appendix. I construct the opinions index by assigning each respondent the left-right orientation score that is typical of someone who answers the questions in the same way that they do. Specifically, I use the whole sample to calculate the average left/right orientation rating of people who “strongly disagree”, “disagree”, “neither agree nor disagree”, “agree” or “strongly agree” with a given statement. I assign each respondent this score for every question they answer, and take the average across the 8 questions that they answered in a given week. By construction, this index has the same mean as the left/right orientation score (2.5). In principle, it can range from 0 to 6; in practice, it has a smaller range of 1.5 to 3.7, with a standard deviation of 0.36.

#### 2.1.4 Empirical strategy

I first use the survey data to examine whether there are changes in the probability of discussing politics with different groups during the week of Thanksgiving. The regression equation is:

$$D_{ict}^j = \alpha + \beta \text{Thanksgiving}_{ct} + \sum_t W_t + \text{Canada}_c + \mathbf{X}_{icw} \gamma + \epsilon_{icw} \quad (1)$$

In this equation  $D_{ict}^j$  is an indicator for whether individual  $i$  from country  $c$  answering the survey at time  $t$  discussed politics or current events with group  $j$  in the previous week.. The key independent variable is  $\text{Thanksgiving}_{ct}$ , which takes the value of 1 if the respondent answers the survey within 7 days after Thanksgiving.  $W_t$  are survey week fixed effects and  $\text{Canada}_c$  is a fixed effect for being Canadian.  $\mathbf{X}_{icw}$  is a set of individual level controls which include the demographic variables listed in Table II, as well as every combination of fixed effects for sample (i.e. did the respondent start the survey in the week of October 5th, or November 17th), Canadian, and indicators for having a family to the right of the respondent and family to the left of the respondent.<sup>7</sup>

---

<sup>7</sup>The sample controls are added because there is a shift in the initial ideological positions of respondents

Standard errors are clustered at the country by survey wave level.<sup>8</sup>

One possible threat to identification in this regression is that there was a Canadian election on October 21st, 2019 - one week after Thanksgiving. Clearly, this could cause people to discuss politics more, as well as potentially change their political opinions. As an initial attempt to address this issue, I include both pre-election and post-election linear time trends interacted with “Canadian” in all of my regressions. In order to test whether my results are being driven by the election (or any other source of endogeneity), I also examine whether my results are present among French Canadians, who do not traditionally celebrate Thanksgiving. I show that there are no similar effects among the French in my sample, which implies that my results are unlikely to be driven by the Canadian election.<sup>9</sup>

I also examine whether discussions are more heated when they take place around Thanksgiving. For this exercise, I use the same regression specification as above, but use the dependent variables capturing recent agreement in discussions, as well as enjoyment of the most recent discussion.

Finally, I examine whether talking to a more diverse set of people at Thanksgiving changes people’s political opinions, and whether this reduces polarization. If individuals are influenced by their families, then we should see that people who have family members to the right of them politically move right after Thanksgiving; the opposite should be true for people whose families are to the left of them. I test this hypothesis by running the following regression:

$$Opinion_{ict} = \alpha + \beta Thanksgiving_{ct} + \delta Thanksgiving_{ct} * Diff_i + \Sigma_t W_t + Canada_c + \mathbf{X}_{icw} \gamma + \epsilon_{icw} \quad (2)$$

In this equation,  $Opinion_{ict}$  is the weekly political opinion index for respondent  $i$  in between sample 1 and sample 2, which differs across the American and Canadian sample. The indicators for having a family to the left and to the right of the respondent are included because I later split the sample along these dimensions; I include them here for consistency. They make no difference to the discussion results.

<sup>8</sup>I have confirmed that there is no issue with over-rejection of the null hypothesis when clustering at this level. Specifically, if I remove the weeks of Thanksgiving and randomly assign “Thanksgiving” separately for the Canadian and American sample within the remaining weeks, repeatedly, I get coefficients that are significant at the 5% level 5% of the time.

<sup>9</sup>Additionally, while I do not show this below, the results appear to be slightly stronger in the U.S. sample compared to the Canadian sample. Results are available upon request.

country  $c$  at time  $t$ . The variable  $Diff_i$  is the right/left difference between the respondent and their family. Positive values of  $Diff_i$  indicate that the respondent’s family is to the right of them, while negative values indicate that the respondent’s family is to the left of them. In this version of the regression specification, I also include  $Diff_i$  as a control in the vector  $\mathbf{X}_{icw}$ , and along with all interactions of  $Diff_i$  with fixed effects and controls.

To interpret this regression as capturing the effect of diversity in social interactions, we must assume that there is no other reason that Thanksgiving would influence opinions in a way that pushes people towards their families’ political views. This assumption could be violated, for example, if there was a general trend towards depolarization after Thanksgiving. Perhaps people are happier or more thankful in this week, and this causes the more extreme members of the sample to move towards the center. Because more extreme people are likely to have families who are closer to the center than they are, a general depolarization could produce this type of pattern.

In order to distinguish between the ”social diversity” and ”general depolarization” explanations for my results, I run two additional regressions. First, I separately examine the interaction between Thanksgiving and the two components of  $Diff_i$ : the respondent’s family’s opinions, and the respondent’s own opinions. A general depolarization would predict that all respondents with the same viewpoints moved in the same direction, while the social diversity hypothesis would predict that respondents with the same political viewpoints move in different directions, depending on the viewpoints of their families. As I show below, the latter prediction holds. Secondly, I examine whether there is a differential change in respondents’ opinions associated with the viewpoints of other social contacts: parents, spouses, coworkers, friends, etc. There is no significant increase in the probability of talking to these groups around Thanksgiving, so any change in the direction of these groups would cast suspicion on the social diversity hypothesis. As I show in the results section, people move only in the direction of siblings and extended family. Both of these regressions support the interpretation that the effect of Thanksgiving operates through its impact on ideological diversity in social interactions.

Once I establish that respondents move in the direction of their families after Thanksgiving, I calculate whether this reduces polarization by predicting the average change in political opinions for people who are at different points on the initial political orientation scale. As I show, people who are far to the left tend to have high values of  $Diff_i$ , while the reverse is true for people who are far to the right. This, combined with the results from Equation 2, implies that Thanksgiving tends to reduce polarization.

## 2.2 Canadian electoral reforms

The MTurk survey establishes directly whether people are exposed to more ideologically diverse viewpoints at Thanksgiving, and whether this results in a change in political opinions. However, there are two key weaknesses of this exercise. First, as shown in Table II, the MTurk sample is not representative of the general population. MTurk respondents are younger, more educated, and more left-wing than the general population. If these characteristics are correlated with “persuadability”, it could be that the treatment effects are larger in this sample than in the general population. Secondly, while I can detect changes in reported political opinion within my sample, it is not clear how economically important these changes are, or whether they even reflect actual changes in respondents’ views. It could be that respondents simply *report* different beliefs after being exposed to family, either because these views are seen as more socially acceptable or because they are more salient.

In order to address these concerns, I next use a series of Canadian electoral reforms that affected the timing of elections to examine whether voting is less polarized when elections occur just after Thanksgiving. Because this analysis applies to the entire voting population, it does not suffer from the sample selection issue present in MTurk. It also captures changes in political *behavior*, as opposed to expressed political opinions, which overcomes the reporting issue.

### 2.2.1 Background

Canada maintains a British-style system of parliamentary democracy, in which members of parliament are elected to represent districts (“ridings”), with federal and provincial leaders (“Prime Minister” and “Premiers”, respectively) being selected based on which party holds the most seats in parliament. Prior to the mid-2000’s, election dates were typically set by the ruling party, under the constraint that they had to hold an election within 5 years of the previous one.<sup>10</sup> This allowed the ruling party to manipulate election timing to increase their odds of victory, and resulted in more elections than were strictly necessary - on average, the gap between elections in the 1980’s and 1990’s was around 3.7 years.

In order to address these problems, most provinces and the federal government introduced legislation during the mid-2000’s to move to a fixed-date system, in which elections

---

<sup>10</sup>An exception to this procedure can occur when the ruling party holds a minority government; I discuss the rules in this case below.

were to be held on a specified date 4 years after the previous election. Table IX shows the year of reform and the election days chosen by each jurisdiction. The days chosen for elections typically followed a rule along the lines of “The 1st Tuesday in November”. Since Canadian Thanksgiving also follows this type of rule (it occurs on the 2nd Monday in October), the fixed election legislation also fixed the timing of elections relative to Thanksgiving. As shown in the table, three jurisdictions - Saskatchewan, Newfoundland, and the Federal government - chose fixed dates that set the election between 1-4 weeks after Thanksgiving. Five of the remaining jurisdictions chose dates within 1-3 weeks prior to Thanksgiving, while the remaining jurisdictions either chose dates in the spring (Alberta and British Columbia), or did not introduce fixed date legislation (Nova Scotia).

For several reasons, the fixed date legislation is not followed with perfect compliance. First, the Canadian Supreme Court has ruled that the legislation is constitutionally non-binding. Most governments do appear to comply with the legislation; however, there have been several high-profile cases where majority governments have ignored fixed dates, such as during the federal election of 2011.<sup>11</sup> Secondly, most fixed date election legislation contains clauses allowing for elections to be moved if they coincide with a significant religious holiday or (in the case of provincial legislation) a federal election. The latter provision has been used occasionally, for example to delay the 2015 Saskatchewan election to 2016. Finally, the fixed date legislation does not always apply in the case of a minority government - one in which the ruling party holds the plurality, rather than the majority, of seats in parliament. In this case, the non-ruling parties can force an election at any time by defeating the government on a major bill, which is known as a vote of non-confidence. I eliminate fallen minority governments from my sample below.

## 2.2.2 Data

I gather data on election results - vote shares by party - for all provincial and federal elections from 1980-2019 from Wikipedia.com.<sup>12</sup> I also use this source to gather other election information, such as voter turnout and the number of parties, which I will use in placebo tests. For Nova Scotia, the only province that did not implement a fixed date

---

<sup>11</sup>The fact that these cases were high-profile and controversial speaks to the underlying factor disciplining governments to comply with fixed date elections: public opinion.

<sup>12</sup>I use province or federal-level data, rather than riding level data, for several reasons. First, the shock I use is at the jurisdiction level; I would expect the main effect to occur across jurisdictions, rather than within jurisdictions. Secondly, riding data is typically only available for a smaller number of years, and often only in textual format. Finally, the riding boundaries change substantially from election to election, preventing me from analyzing changes within ridings.

reform, I assign a placebo “treatment” year equal to the median treatment year in the sample (2011).<sup>13</sup> The final sample is a set of 110 elections, which occur between 1980-2019.

If individuals influence each other at Thanksgiving, there should be movement towards each jurisdiction’s political majority around this time. In this case, I would expect to see that there is more agreement among the electorate when elections are held immediately after Thanksgiving. I measure agreement in voting through indices that capture equality in vote shares. The underlying intuition is that equality in vote shares across parties corresponds to the maximum level of disagreement within the population. Conversely, perfect inequality in vote shares - all votes going to the same party - captures the maximum amount of agreement. I therefore use measures that are common in the literature on income inequality. The specific measures I use are i) the maximum vote share, ii) the standard deviation of vote shares across parties, iii) the Gini index of vote shares across parties, and iv) a generalized entropy index.<sup>14</sup> All four variables increase when there is more inequality in vote shares, or more agreement among the electorate.

I also include controls for demographics in my regressions. These variables come from the annual Canadian Labour Force Survey (LFS), which provides monthly data on employment and labor market outcomes for the population aged 15 and older. For each jurisdiction and election year, I use the LFS to calculate i) the proportion of the adult population that was aged 15-24, or 55 plus, ii) the proportion of the adult population that had a college degree, and iii) the proportion of the adult population that was female, and iv) the total population aged 15 and older.<sup>15</sup>

### 2.2.3 Empirical strategy

My empirical strategy is to compare electoral outcomes across jurisdictions that set fixed dates either after Thanksgiving or at some other point during the year, before and after the fixed date reforms.<sup>16</sup> My main regression equation is:

---

<sup>13</sup>Results are not sensitive to assigning different treatment years for Nova Scotia, or to excluding Nova Scotia entirely.

<sup>14</sup>The generalized entropy index is calculated as  $\sum_j \frac{p_j \log(p_j)}{\log(N)}$ , where  $j$  indexes parties in an election,  $p_j$  is the vote share of party  $j$ , and  $N$  is the number of parties. This index is equal to 0 when a single party holds all the votes, and is equal to -1 when vote shares are evenly distributed across parties.

<sup>15</sup>For elections in my sample occurring prior to 1990, I interpolate values using Census information from 1981 and 1986.

<sup>16</sup>In principle, I have three groups: those that set elections just after Thanksgiving, those that set elections just before Thanksgiving, and those that held elections at other times of the year. In practice, the latter two groups look identical; as a result, I simply compare the post-Thanksgiving group to the other two groups.

$$V_{jt} = \alpha + \beta Treatment_j * Post_t + \gamma Post_t + \Sigma_j P_j + \delta \mathbf{X}_{jt} + \epsilon_{jt} \quad (3)$$

where  $V_{jt}$  is a measure of outcomes for an election taking place in jurisdiction  $j$  at time  $t$ ;  $Treatment_j$  is an indicator for whether jurisdiction  $j$  is one of the three jurisdictions that chose post-Thanksgiving election dates;  $Post_t$  is an indicator for whether the election occurred after fixed-date election reform;  $P_j$  are jurisdiction fixed effects; and  $\mathbf{X}_{jt}$  is a vector of control variables that includes demographics, day of week fixed effects, and a quadratic time trend. Standard errors are clustered at the jurisdiction level.

The identification assumption in this regression is that there is no other reason for electoral outcomes to change in the treatment group, relative to the control group, after the reforms. One particular violation of this assumption would occur if the two sets of jurisdictions had different trends prior to the reforms. I show visually that this is not the case, using event study graphs. Alternatively, it could be the case that other policy changes or shifts in political attitudes took place in the treatment jurisdictions around the time of the reform. I show, however, that there are no changes in other election policy variables, or in the vote shares of specific parties consistently across the three jurisdictions.

## 3 Results

### 3.1 Canada-U.S. comparison

#### 3.1.1 Changes in exposure to different opinions

First, I examine how Thanksgiving affects the diversity of interactions, as measured through political discussions. Table IV shows how Thanksgiving affects the probability of discussion with different groups of people. The table shows that Thanksgiving has a large and significant impact on discussion with two groups: siblings and extended family. During the week of Thanksgiving, respondents are 9.6 percentage points more likely to say that they had discussed politics or current events with siblings. Recall from Table III that the average probability across all weeks is around 19 percentage points; this therefore represents a 50% increase over the mean. For extended family, the increase is 7.8 percentage points, approximately a 60% increase from the mean of 13.3%. Figure I shows the combined effect on the probability of talking to either a sibling or extended family member, separately by week. There is a sharp increase in the week immediately after Thanksgiving, which falls to about half its size in the following week, before returning

to the pre-Thanksgiving level in the following week.<sup>17</sup> Returning to Table IV, there are smaller, but non-significant, increases in the probability of talking to parents, adult children, friends, and other groups (all around 4.5 percentage points), and a non-significant 2.5 percentage point decrease in the probability of talking to a spouse or partner.

How does this change in discussion rates affect the diversity of political opinions that MTurk respondents are exposed to? Table V shows how the level of agreement in political discussions varies during Thanksgiving. In the first column, the dependent variable is the average absolute difference in left/right orientation between a respondent and his or her discussion partners in a given week. This increases over Thanksgiving, but only slightly - an increase of 0.041 on a base of about 1.1 - and the results are not significant. This likely occurs because individuals appear to continue to have their normal level of discussion with other groups; the increase in discussion with siblings and extended family appears is in addition to a slightly higher rate of discussion with other people, which means that the effect on the average will be somewhat limited. Instead, it seems likely that this change increases the probability of encountering views that are further away from your own. Column (2) confirms this intuition. In this column, I show the change in the *maximal* disagreement between a respondent and their discussion partners from the previous week. This variable increases by 0.125, compared to a mean of around 1.6, and the difference is statistically significant at the 1% level.

In columns (3)-(5), I look instead at reported levels of agreement and/or enjoyment in the respondent's most recent political discussion. Respondents report a slightly lower level of agreement on a scale of 1-10, are less likely to report that they enjoyed the discussion a lot, and more likely to report that they did not enjoy the discussion at all. While only the "enjoyed a lot" variable is marginally significant, note that the effects on the latter two variables are large: the probability of enjoying a discussion a lot falls by 6.7 percentage points, compared to a mean of 28.5%, and the probability of not enjoying the discussion increases by 4.4 percentage points on a base of 9.4 percentage points.

In sum, Thanksgiving appears to increase the probability that we discuss politics with our siblings and extended family, groups that we disagree with relatively more than the people we usually talk to. For respondents in my sample, this increases the range of political orientations to which they are exposed, and reduces their enjoyment during the

---

<sup>17</sup>The omitted category in all of my graphs is the combination of "more than 21 days pre-Thanksgiving" and "more than 28 days post-Thanksgiving". This is because the former category is made up entirely of Americans, and the latter category is made up entirely of Canadians. If I did not combine them, the coefficients on either the "-3" or "5+" variables would be biased, because they would reflect constant difference between the two countries.



discussions. Next, I ask: is this exposure to a more diverse set of opinions effective in reducing polarization?

### 3.1.2 Changes in opinions and polarization

Table VI examines how my respondents' weekly left/right opinion index changes in the week after Thanksgiving, and how this varies with the political orientation of a respondent's siblings and/or extended family members. Recall from Figure II that the typical MTurk respondent in my sample is fairly left-wing, while their siblings and extended families are more moderate. If families are effective in changing respondent's political opinions, then, I should see a rightward shift in MTurk respondent's opinions. The first column of Table VI confirms that the sample moves slightly to the right, although the change is small and not statistically significant. The opinion index is 0.016 points higher in the week of Thanksgiving, compared to a mean of about 2.5. Of course, the change in the index should depend on whether a specific respondent's family is to the left or to the right of them. In columns (2) and (3), I split the sample into those whose families are (weakly) to the right of them, and the sample whose families are to the left of them.<sup>18</sup> People whose siblings and extended families are at the same or a higher level on the left/right orientation index move significantly right during Thanksgiving, while people whose families are to the left of them move left. The coefficients for the two samples are of a similar magnitude, about 0.035 in absolute value. The fourth column shows that the difference between the two samples is (marginally) statistically significant.

Of course, these coefficients only captures the effect during the week after Thanksgiving. How long does the effect last? Figure III and Figure IV show how the effect changes over time. For the sample with families to the right, the effect is primarily concentrated in the week after Thanksgiving, although the coefficients remain above zero for up to 4 weeks afterward. For the sample with families to the left, the effect remains relatively constant for the 4 weeks following Thanksgiving. On net, then, the effect seems to last between 1-4 weeks.

In column (5) of Table VI, I examine how the change in opinions varies when the difference between a respondent and their family grows. I interact "Thanksgiving" with

---

<sup>18</sup>Those with families at the same left/right orientation as them also move slightly to the right during the week of Thanksgiving, which is why I group them with the "family right" sample. This could be due to reporting error in the left/right orientation index, or to a general rightward shock in the sample. As I will show in the next columns of the table, the degree to which a respondent moves right increases as their family moves to further to their right, which is consistent with the story that families are influencing political opinions.

“L/R family - L/R self”. The latter variable is positive when a respondent’s family is to their right and negative when a respondent’s family is to their left; it gets larger in absolute value when there is a bigger difference in orientation between the respondent and their family. The coefficient on the interaction term is positive and statistically significant, indicating that respondents whose families are further to the right show a bigger rightward change in their opinion index.

If these changes are truly driven by social interactions with family, and not by a general depolarization in my sample, the effects in column (5) should be driven by the “L/R family” variable, rather than the “L/R self” variable. In other words, we should see that people with identical initial political orientations move in different directions depending on the orientation of their siblings and extended families. Column (6) of Table VI shows that this is the case. The coefficient on Thanksgiving interacted with family orientation is positive and significant, while there is no significant change associated with whether the individual was initially more left or right wing.<sup>19</sup> Column (7) provides further evidence that the effect I document works through interaction with siblings and extended family, by examining whether there are similar changes associated with the political beliefs of other groups: other family (spouse, parents, and adult children), or non-family (coworkers, friends, and neighbors). Because there is no increase in the probability of interacting with these groups around Thanksgiving, significant coefficients on these variables would cast doubt about the channel driving my results. However, the results show that the effect is confined to the groups that show a first-stage response.

How do these changes affect polarization in the sample? The typical left-wing member of my sample has family that is to the right of them. Respondents who answered that they were below 2 on the initial left/right orientation scale have families that are, on average, 1.6 points to their right. Using the coefficients in Table VI, this predicts that these respondents move 0.027 points to the right on the opinion index. Conversely, the right-wing members of my sample typically have families that are to their left. Respondents who answered that they were above a 4 on the initial left/right orientation scale have families that are 0.9 points to their left, on average. The predicted change in their left/right opinion index is therefore -0.016. The left and right wing members of my sample therefore move closer together on the opinion index, by about 0.045 points, or around 12% of a standard deviation.

It is important to note that 12% of a standard deviation is a lower bound on the effect

---

<sup>19</sup>Results are similar if I include interactions between Thanksgiving and indicators for each possible level of a respondent’s initial political orientation.

size of talking to family, since I have not yet scaled it by any first stage effect. While I could do this in principle, using the effect on political discussions documented in the last subsection, I am reluctant to do this for two reasons. First, it is unclear whether the treatment that leads to reduced polarization is “discussing politics with family”, or simply “seeing family”. It is possible that socializing with people who are different from you politically can influence your opinions, even without political discussion. Since I do not capture this effect in my discussion variables, the true first-stage may be larger than what I have documented. Secondly, there may be multiplier effects if people first discuss politics with family, and then with friends or other people within the same week. In the week of Thanksgiving, I may get exposed not only to my own family’s opinions, but also indirectly to the opinions of the family of my friends or coworkers. Again, this will imply that I have under-measured the size of the shock in the first stage. If I ignore these concerns and scale by the first stage estimate (a 12% increase in the probability of talking to siblings or extended family members), the implied effect size is around 1 standard deviation. This is probably too large, but is useful for providing an upper bound on the effect size. I conclude that the effect of widening your social circle at Thanksgiving is somewhere between a 0.12 and a 1 standard deviation reduction in polarization. For comparison, Hunt Allcott and Gentzkow (2020) document that inducing Facebook users to deactivate their accounts for 4 weeks, with near perfect compliance, reduces polarization by around 16% of a standard deviation. The effect of Thanksgiving is therefore approximately equal to or larger than that of deactivating Facebook for 4 weeks.

### 3.1.3 Ruling out threats to identification

There are two key threats to identification in my exercise above. First, because I do not observe every member of my sample in every week, it may be the case that an observably different subset of respondents answer my survey in the week just after Thanksgiving. If this is the case, my results could be driven by having a sample that was already less polarized prior to Thanksgiving. Second, as mentioned previously, there was a Canadian election on October 21st - one week after Canadian Thanksgiving. If people become less polarized prior to elections, then it could be the election that is driving my results.

Table VII attempts to alleviate the first of these concerns. In columns (1)-(3), I first check whether there is any difference in the *initial* L/R opinion index between the set of respondents who answer during the week of Thanksgiving and other respondents. The dependent variable is a respondent’s opinion index in the first week that they were in

the sample. I assign this index to them for every week, and check whether the average value of this variable changes during Thanksgiving. If so, this is an indicator that the people who select out of my sample during the week of Thanksgiving are different, on average, then respondents who remain in the sample. Column (1) shows that there is no overall difference in the initial opinion index between respondents who answer the week of Thanksgiving and other respondents. Columns (2) and (3) repeat this exercise for the sample with families to the left of them and families to the right of them, to mirror the results in Table VI. There is also no difference in initial opinions for this group.

A second way of checking whether sample selection is driving my results is to use individual fixed effects in my regressions, which compare each respondent’s opinion index in the week of Thanksgiving to their own average in weeks outside of Thanksgiving. This helps with the selection problem because the estimator only uses people who are observed in both Thanksgiving and non-Thanksgiving weeks for identification. Columns (4) and (5) of Table VII show the results of fixed effects regressions, for the sample with family to the right and to the left of them separately. The results are almost identical in size to the coefficients in Table VI, and are significant at the 1% and 5% level for the “family right” and “family left” samples respectively. Again, this suggests that non-random attrition is not driving my results.

In Table VIII, I perform a placebo test using the French Canadian sample. French Canadians do not traditionally celebrate Thanksgiving, and should therefore be unaffected by Thanksgiving. They are, however, affected by the Canadian federal election. If the election was driving my results, I would expect to see similar results in French Canada as in the rest of the sample. I first assign each respondent a probability of being French based on i) the language they chose to take the survey in, and ii) their geographic location.<sup>20</sup> I interact this variable with “Thanksgiving” and repeat my regressions on family discussions and opinions. The results show that there is no effect in the French sample, either for family discussions or for opinions. While the difference between the French and non-French sample is not statistically significant (except for marginally in the case of family discussions), the point estimates are either close to zero (for the family discussion variable) or small and in the wrong direction (for the opinion variables.) These results

---

<sup>20</sup>Specifically, this variable is equal to 1 if the respondent took the survey in French; 0.7995 if the respondent answered the survey in English but is in Quebec; and 0.2835 if the respondent answered the survey in English but is in New Brunswick. The proportions in the latter two cases are based on the proportion of French speakers in these provinces from the 2011 Canadian Census. The proportion French in other areas of Canada is less than 2%. Unfortunately, I did not ask respondents directly whether they were French, which is why I use this proxy.

are inconsistent with the Canadian election driving my results.

## 3.2 Canadian electoral reform

The previous exercise showed that Thanksgiving reduces opinion polarization among MTurk respondents. It is difficult to interpret these results, however, for two reasons. First, the MTurk sample is not representative of the general population. If young, educated, and liberal people are more easily persuaded, then the effects might be limited to this group. Secondly, it is difficult to say whether the changes in the opinion index are economically meaningful. Do they correspond to changes in political behavior?

To answer these questions, I next examine the impact of holding an election after Thanksgiving, using the Canadian electoral reforms described in previous sections. I compare measures of electoral agreement in jurisdictions that set their fixed date elections just after Thanksgiving (the “treatment” group) to other jurisdictions (the “control” group), before and after the reforms. For this comparison to be valid, it must be the case that there is no other reason for voting behavior to change in the treatment group, relative to the control group. In support of this assumption, Table X provides evidence that these jurisdictions looked very similar prior to the reform. In this table, I compare characteristics of elections between the treatment and control groups, using elections that happened prior to the reforms. The table shows that the main difference between these groups was the timing of elections: the treatment group was already more likely to hold their elections in the month after Thanksgiving prior to the reforms. As I will show, however, there was a large increase in this probability after the reforms. The treatment jurisdictions were also slightly younger than the rest of the country, and have a larger population because the Federal government is included in the treatment group. There was no significant difference in voting polarization, general political orientation, voter turnout, or other demographic characteristics of the electorate prior to the reform. Of course, the difference-in-difference estimate will eliminate any fixed differences between the treatment and control group; the key concern is whether there are any *changes* in observed or unobserved characteristics that might explain my results. I provide further evidence below that this is not the case.

### 3.2.1 Election timing and voter agreement

Table XI shows that the electoral reforms had a significant impact on the timing of elections. As shown in Column (1), the probability of holding an election within 4 weeks

after Thanksgiving increased by 48 percentage points in the treatment group, compared to the control group. As shown in Column (2), this corresponds to a decrease of about 150 days since the previous Thanksgiving. Figure V shows that this increase occurs sharply at the time of the reforms. If Thanksgiving does result in less polarization, we should therefore see more agreement among the jurisdictions that became more likely to vote in the month after Thanksgiving.

Table XII shows how measures of agreement in voting changed in the treatment group, relative to the control group. In the first column, I examine a simple measure of agreement among the voting population: the maximum vote share across parties. As shown in the table, the maximum vote share increases by around 13 percentage points in the treatment group compared to the control group after the reforms. The upper left panel of Figure VI shows that this effect is not driven by different trends between the two groups prior to the reform. The effect occurs sharply at the time of reform, and remains persistent up to 3 elections following the reform. The next columns of Table XII use alternative measures of inequality in voting which take into account the entire distribution of votes across parties. These results show that there are significant increases in all measures of vote share inequality. The standard deviation of votes increases by 0.059, the Gini index by 0.098, and the entropy index by 0.112 in the treatment group following the reform. The remaining panels of Figure VI show the results for these variables visually. In all cases, there is no differential trend prior to polarization, with increases in the indices immediately following reform.

How large are these effects? A 13 percentage point increase in the maximum vote share is clearly quite substantial, especially given that the “first stage” effect of the reforms on the timing of elections is not perfect.<sup>21</sup> For comparison, Martin and Yurukoglu (2017) show that watching Fox News for 1 hour per week for a year increases the Republican vote share by around 7 percentage points. The increases in the other variables are equivalent to 15-30% of their pre-reform means, or around 1 standard deviation in each case. The effect of holding elections immediately after Thanksgiving appears to be very large.

---

<sup>21</sup>As in the previous exercise, it is not clear that we should scale the reduced form effects of the reforms by the first stage effect of holding an election within 4 weeks of Thanksgiving. For one thing, the reforms affect the timing of elections within the month after Thanksgiving as well, which could affect voting in a non-linear way. Secondly, the effect of the reforms on election timing appeared to be strongest in the first election following the reform. If there are incumbency effects, than some of the continued effect on voting could be due to this.

### 3.2.2 Ruling out threats to identification

One of the key limitations of this analysis is the limited number of observations. With just 3 jurisdictions in my treatment group, it is possible that my results could be driven by random events in a single jurisdiction. Is this the case? In table Table XIII, I break down the estimates separately for the three treatment jurisdictions.<sup>22</sup> All three jurisdictions show significant changes in polarization. Saskatchewan and Newfoundland saw large changes in the maximum vote share, while the Federal government saw significant changes at other points in the distribution. The fact that I see similar results in all three jurisdictions suggests that my results are not arising by chance.

More generally, I would like to rule out the possibility that my results are driven by any other change, legislative or otherwise, that occurred around the same time as fixed date legislation and disproportionately affected the treatment group. While I cannot eliminate this possibility entirely, I show in Table XIV that there were no other observable changes in the characteristics of elections in the treatment group compared to the control jurisdictions. In columns (1)-(2), I examine two variables that proxy for changes to election policy: the number of seats per capita, and the years since the previous election. Neither of these variables showed a significant change in the treatment group, relative to the controls.

In the next two columns, I examine whether my results are driven by changes in either the number of political parties, which could affect my dependent variables mechanically, or by voter turnout. There is no change in the number of parties in the treatment group compared to the control. There is a marginally significant reduction in voter turnout, of around 8 percentage points. I investigate this result visually in Figure VII. The reduction in turnout in the treatment group compared to the control occurs largely before the reforms, although there is perhaps a small change post-reform. It seems unlikely that these patterns can explain the sharp discontinuities in voter behavior shown in Figure VI.

Finally, in columns (5)-(7), I examine the vote shares of the three major parties in Canada, which is a proxy for changes in the political attitudes of Canadians more broadly. There were no changes in any these variables, suggesting that the treatment group did not move to the left or to the right relative to the rest of the country.

In sum, the reforms did not coincide with significant changes in any other facet of

---

<sup>22</sup>Technically, the results could also be driven by an outlier in the control group. This is not the case: my main results result from increases in the treatment group, with no change in pre/post behavior in the control group.

election policy or political behavior in Canada. The only significant change in the post-reform period was that the treatment jurisdictions saw substantially greater agreement among their electorate, compared to the pre-reform era and to the control group.

## 4 Conclusion

Both the MTurk and electoral reform results show that talking to people we disagree with at Thanksgiving reduces political polarization. The effects are quite large compared to other factors that have received more attention in the literature, such as social media or cable news, and suggest that reduced ideological diversity in social interactions may be an important part of the explanation for increased polarization.

There are two potential policy implications of my results. First, more narrowly, they suggest that our families may be a valuable and under-recognized resource when it comes to reducing political polarization. Many authors have pointed out the role of identity in driving polarization, especially as political orientation comes to increasingly overlap with demographic characteristics such as education or urban status (e.g., Klein, 2020). One way to combat this tendency is to use the fact that we have multiple identities, such as family membership, that may be less correlated with our politics.

More broadly, the results suggest that talking to people who disagree with us can, under the right conditions, reduce political disagreement. How might we encourage this kind of behavior? While my results cannot speak to the effect of any specific policy, it is worth pointing out that the “shock” in my analysis - Thanksgiving - is simply a period of coordinated leisure. As with any other behavior, reducing the costs of socializing is likely to be effective in increasing its frequency.

Author affiliation: Kirsten Cornelson is an assistant professor at the University of Notre Dame.

## 5 Appendix

### 5.1 Construction of the political opinions index

To evaluate respondents’ left-right positions on a weekly basis, I ask respondents to report their level of agreement with a set of 8 statements that are randomly chosen each week from a pool of 20. The statements were designed to mimic the left-right ratings from



the Comparative Manifesto Project (Krause et al., 2020), which rates political parties from around the world based on statements from their platforms. The CMP rates parties as more right-wing if they: express support for the military, emphasize personal freedom, express support for a status quo interpretation of the constitution, emphasize the importance of political authority, express support for free markets and economic orthodoxy, believe in limiting the welfare state, emphasize law and order, and express support for a national way of life and traditional morality. It rates parties as more left-wing if they: express anti-imperialist sentiments, make negative statements about the military or military spending, explicitly make pro-peace statements, promote international cooperation, support market regulation and economic planning, support an expansion of the welfare state, make statements that are pro-organized labour, and explicitly emphasize democratic institutions.

Based on these principles, I developed a series of 24 statements and asked a group of MTurk workers to rate the political beliefs of someone who strongly agreed with each statement. All statements were associated with people on the correct side of the political scale (i.e., the statements meant to capture left-wing beliefs were always associated with more left-wing ratings.) In 4 cases, however, this association was not large or statistically significant; I dropped these statements. 3 of these statements were intended to be right-wing, while 1 was intended to be left-wing. The remaining 20 statements were:

#### **Right wing statements**

1. A strong military is important to keep this country safe.
2. Constraints on government are necessary to protect individual liberties.
3. It is the duty of judges and lawmakers to uphold the constitution as written, without imposing their own interpretation of the law.
4. Free markets do more to create prosperity than the government ever could.
5. My family has to live within its means; the government should too.
6. Spending on social programs in this country has gotten out of control. Individuals should not expect the government to provide for all of their needs.
7. I am proud of my country, and anyone who doesn't feel the same way should leave.
8. Traditional family values are under threat in this country.

9. The courts are too easy on criminals these days.

### **Left wing statements**

1. The military budget would be better spent on social programs like health or education.
2. The government should make every effort to resolve disputes peacefully before resorting to military intervention.
3. Our country has a responsibility to provide aid to poorer countries.
4. Our country should not interfere in the politics of other nations in order to advance our own interests.
5. Government regulation is necessary to ensure that markets produce fair outcomes.
6. The government should ensure that basic necessities like food and water do not become too expensive.
7. The government should nationalize vital industries to ensure this country's security and shared prosperity.
8. We need subsidized childcare in this country.
9. A college education should be free to everyone.
10. Strong unions are vital to protect workers' interests.
11. A free press plays a vital role in maintaining our democracy.

To construct the weekly political opinions index, I regressed respondents' political *orientation* index (the self-report of left-right orientation taken in the baseline survey) on their responses to each question, where each level of agreement from 0 to 4 was treated as a categorical variable. Based on this regression, I predicted a level of left-right orientation for each respondent based on their response to each question in each week. Their weekly political opinion index is the average of the predictions based on the 8 statements they evaluated.

## References

- Abramowitz, A. I. and K. L. Saunders (2008). Is polarization a myth? *The Journal of Politics* 70(2), 542–555.
- Boxell, L., M. Gentzkow, and J. M. Shapiro (2020). Cross-country trends in affective polarization. *Working paper*.
- DellaVigna, S. and E. Kaplan (2007). The fox news effect: media bias and voting. *Quarterly Journal of Economics* 122(3), 1187–1234.
- Gentzkow, M. (2016). Polarization in 2016. *White paper*.
- Gentzkow, M. and J. M. Shapiro (2011). Ideological segregation online and offline. *Quarterly Journal of Economics* 126, 1799–1839.
- Hunt Allcott, Luca Braghieri, S. E. and M. Gentzkow (2020). The welfare effects of social media. *American Economic Review*, 629–676.
- Kevins, A. and S. N. Soroka (2018). Growing apart? partisan sorting in canada, 1992-2015. *Canadian Journal of Political Science/Revue canadienne de science politique* 51(1), 103–133.
- Klein, E. (2020). *Why We’re Polarized*. Profile.
- Krause, W., P. Lehmann, J. Lewandowski, T. Matthieb, N. Merz, S. Regal, and A. Werner (2020). Comparative manifesto project, version 2020-5.
- Levy, R. (2020). Social media, news consumption, and polarization: evidence from a field experiment. *Working paper*.
- Martin, G. J. and A. Yurukoglu (2017). Bias in cable news: persuasion and polarization. *American Economic Review* 107(9), 2565–2599.
- Sunstein, C. R. (2017). *Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Taibbi, M. (2019). *Hate, Inc.: Why Today’s Media Makes us Despise One Another*. OR Books.

## 6 Tables and Figures

Table I. Changes in time spent with different groups, 1965-2018

	Average minutes per day	
	1965	2010-2018
Spouse	170	173
Spouse (conditional)	219	342
Children	161	130
Children (conditional)	256	338
Other household members	21	47
Other household members (conditional)	177	262
Non-HH relatives and friends	100	100
Coworkers	135	183
Coworkers, outside of work*	20	13
Neighbors	29	7
Other	65	36

This table shows summary statistics on the number of minutes per day spent in the company of different groups, using time use surveys from 1965 and 2010-2018. The 1965 data come from the Americans' Use of Time Survey, and were provided by ICPSR. The 2010-2018 data are from the American Time Use Survey, and were collected from IPUMS Time Use.

\* This calculation is based on microdata from IPUMS American Heritage Time Use Extract Builder. Data in the second column are from 2010-2012, rather than 2010-2018.

Table II. Summary statistics, MTurk sample

	Both	Canadian	American
Canadian	37.3%	100.0%	0.0%
Age - 18 to 25	14.9%	21.8%	10.8%
Age - 26 to 45	64.4%	66.4%	63.2%
Age - 46 and older	20.7%	11.8%	26.0%
Male	58.0%	60.3%	56.6%
White	75.2%	69.5%	78.6%
Black	7.2%	4.3%	8.9%
Asian	10.3%	16.4%	6.8%
High school or less	9.9%	9.2%	10.3%
Some college	23.4%	19.5%	25.6%
College degree	49.8%	53.7%	47.5%
Post-graduate degree	16.1%	15.8%	16.2%
Left/right orientation (0-6)	2.55	2.49	2.59
Observations	933	348	585

This table shows means of demographic variables and political orientation in the MTurk sample. The left-right orientation variable comes from direct reports from respondents in the first survey wave, where answers range from “extremely left wing” (0) to “extremely right wing” (6).

Table III. Discussion rates and level of disagreement with different groups

	Discussed in past week	L/R distance
Anyone	86.0%	
Friends	46.8%	1.029
Spouse/partner	45.0%	0.802
Parents	33.5%	1.449
Coworkers/other	25.1%	1.610
Siblings	19.4%	1.363
Ext. family	13.3%	1.766
Adult children	7.9%	1.187

The first column of this table shows the average percentage of the MTurk sample that had a political discussion with the indicated group in the past week, over all respondents in the sample. The second column shows the average absolute difference between a respondent’s own 0-6 left/right orientation, and the reported left/right orientation of the group (both reported in the respondent’s first survey wave), for the sample that gave a left/right rating for the indicated group. The left/right distance measure for coworkers/other is an average of the respondent’s ratings of the left/right orientation of their coworkers, and the people in their neighborhood.

Table IV. The effect of Thanksgiving on political discussions

	Any discussion with:						
	Spouse	Parents	Siblings	Adult children	Ext. family	Friends	Other
Thanksgiving	-0.0291 (0.0215)	0.0467 (0.0425)	0.0964** (0.0392)	0.0210 (0.0130)	0.0783*** (0.0269)	0.0412 (0.0256)	0.0434 (0.0272)
Observations	3,859	3,859	3,859	3,859	3,859	3,859	3,859
R-squared	0.073	0.056	0.027	0.127	0.034	0.034	0.047

This table shows the results from regressions of variables indicating that the respondent had a political discussion with the indicated group in the past week on an indicator for “Thanksgiving”, fixed effects, and covariates. Fixed effects include wave fixed effects, an indicator for “Canadian”, and all interactions of Canadian, sample (waves 1-6 or 7-12), and indicators for whether the respondent is to the left or right of their family. Covariates include controls for age, sex, race, education, indicators for each level of an initial rating of the respondent’s left/right position, and the absolute number of days before/after the Canadian election interacted with “Canadian”. Standard errors are clustered at the survey wave by country level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.10

Table V. The effect of Thanksgiving on the level of disagreement in most recent discussion

	Avg. L/R disagreement	Max L/R disagreement	Agreement (0-10)	Enjoyed a lot a lot	Did not enjoy
Thanksgiving	0.0419 (0.0382)	0.123*** (0.0379)	-0.0321 (0.238)	-0.0665* (0.0385)	0.0442 (0.0272)
Observations	3,263	3,263	3,225	3,225	3,225
R-squared	0.072	0.080	0.034	0.060	0.019

This table shows the results from regressions of variables indicating the level of agreement in a respondent's political discussions on an indicator for "Thanksgiving", fixed effects, and covariates. The dependent variable in the first column is the average absolute difference in left/right orientation between the respondent and any groups that he or she had a discussion with in the past week. The dependent variable in the second column is the maximum absolute difference in left/right orientation between the respondent and any person/group she had a discussion with in the past week. The dependent variables in the next three columns are indicators of agreement and enjoyment that are reported by the respondent with reference to their most recent political discussion. In all columns, the sample is the set of individuals who had any political discussion in the past week; in columns (3)-(5), I additionally impose the restriction that the individual had non-missing responses to the agreement and enjoyment questions in that week. Fixed effects include wave fixed effects, an indicator for "Canadian", and all interactions of Canadian, sample (waves 1-6 or 7-12), and indicators for whether the respondent is to the left or right of their family. Covariates include controls for age, sex, race, education, indicators for each level of an initial rating of the respondent's left/right position, and the absolute number of days before/after the Canadian election interacted with "Canadian". Standard errors are clustered at the survey wave by country level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.10

Table VI. The effect of Thanksgiving on political opinions

	L/R opinion	L/R opinion	L/R opinion	L/R opinion	L/R opinion	L/R opinion	L/R opinion
Thanksgiving	0.0161 (0.0190)	0.0330** (0.0159)	-0.0381 (0.0415)	-0.0381 (0.0412)	0.00170 (0.0174)	-0.0406 (0.0310)	0.012 (0.015)
Thanksgiving x Family right				0.0711* (0.0402)			
Thanksgiving x Difference					0.0171** (0.00740)		0.034*** (0.011)
Thanksgiving x family L/R index						0.026** (0.0106)	
Thanksgiving x self L/R index						-0.0110 (0.0069)	
Thanksgiving x Difference (oth. family)							-0.028* (0.015)
Thanksgiving x Difference (non- family)							0.004 (0.009)
Sample	All	Family right	Family left	All	All	All	All
Observations	3,859	2,830	1,029	3,859	3,859	3,859	3,762
R-squared	0.441	0.414	0.293	0.465	0.471	0.471	0.512

The first column of this table shows the results from regressions of a respondent's weekly left-right opinion index on an indicator for "Thanksgiving", fixed effects, and covariates. The second and third columns show the same regression, but where sample is split into respondents whose families are to the right of them politically, and those whose families are to the left of them politically. Column (4) tests whether the difference between the coefficients in columns (2) and (3) is statistically significant. Column (5) shows the results when "Thanksgiving" is interacted with a variable measuring the difference between the respondent and their family on the left-right orientation scale. Higher values of this difference indicate that the family is further to the right of the respondent. Column(6) breaks the difference in column (5) up into its component parts, separately examining the relationship between changes in political opinion related to family political position and own initial political orientation. Column (7) shows the results when interaction terms are added between Thanksgiving and the political difference between the respondent and other groups: "other family" (spouse, parents, and adult children), and "non-family" (friends, neighbors, and coworkers). The sample is slightly smaller in this case because a respondent must have non-missing values on more variables. Fixed effects and controls are the same as in previous tables, except that in columns (4) and (5), the fixed effects and covariates are interacted with "Family right" and "(L/R family-L/R self", respectively. Standard errors are clustered at the survey wave by country level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.10



Table VII. Ruling out selection as a mechanism

	Dependent variable:					
	L/R opinion (first wave)	L/R opinion (first wave)	L/R opinion (first wave)	L/R opinion (first wave)	L/R opinion	L/R opinion
Thanksgiving	0.00176 (0.0222)	-0.00208 (0.0197)	0.00717 (0.0523)	-0.00410 (0.0203)	0.0363*** (0.00878)	-0.0359** (0.0153)
Thanksgiving x (L/R family - L/R self)				0.00297 (0.00840)		
Sample	All	Family right	Family left	All	Family right	Family left
Individual FE					X	X
Observations	3,859	2,830	1,029	3,859	2,830	1,029
R-squared	0.442	0.412	0.361	0.477	0.831	0.778

The first column of this table show the results from regressions of a respondent's *initial* left/right opinion index, taken in their first survey wave, on an indicator for "Thanksgiving", fixed effects, and covariates. The second and third columns show the same regression, but where sample is split into respondents whose families are to the right of them politically, and those whose families are to the left of them politically. The fourth column uses the full sample, but adds an interaction between "Thanksgiving" and the difference between the family's left/right position, and the respondent's left/right position. The fifth and sixth columns of this table show that the main results (using the *weekly* left/right opinion index) are robust to adding individual fixed effects to the regression. All other regression details are the same as in the main regression tables. \*\*\* p<0.01, \*\* p<0.05, \* p<0.10

Table VIII. Placebo test using French speakers

	(1) Discussion w. family	(2) L/R opinion	(3) L/R opinion
Thanksgiving	0.128*** (0.0423)	0.0387** (0.0161)	-0.0595** (0.0275)
Thanksgiving x French	-0.109* (0.0621)	-0.110 (0.0644)	0.150 (0.117)
Sample	All	Family right	Family left
Observations	3,859	2,830	1,029
R-squared	0.029	0.417	0.295

This table shows that there are no treatment effects of Thanksgiving among French Canadians (who do not celebrate Thanksgiving). The variable “French” is a proxy for French-speaking status which is equal to 1 if the respondent answered the survey in French; 0.7995 if the respondent answered the survey in English but is in Quebec; and 0.2835 if the respondent answered the survey in English but is in New Brunswick. The proportions in the latter two cases are based on the proportion of French speakers in these provinces from the 2011 Canadian Census. All other regression details are the same as in the previous tables, except that the variable “French” has been added to the controls. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$

Table IX. Dates of Canadian electoral reforms

Jurisdiction	Law passed	Fixed date (every 4 years)
British Columbia*	2001	2nd Tuesday in May
Newfoundland	2004	2nd Tuesday in October
Ontario	2005	1st Thursday in October
Federal	2007	3rd Monday in October
New Brunswick	2007	4th Monday in September
PEI	2007	1st Monday in October
Saskatchewan**	2007	First Monday in November
Manitoba	2008	1st Tuesday in October
Alberta	2011	Between March and May
Quebec	2013	1st Monday in October
Nova Scotia	NA	NA

This table shows the year in which each jurisdiction undertook its fixed date electoral reform in Canada, and the day on which the election is now supposed to occur.

\*Moved to 3rd Saturday in October in 2017; no elections since change

\*\*Moved to last Monday in October in 2018; no elections since change

Table X. Summary statistics and balancing test for Canadian elections (pre-reform)

	Election date after reform:		
	Post-Thanksgiving	Other	Difference
<i>Election dates:</i>			
Within 4 weeks of Thanksgiving	0.250	0.069	0.181**
Days since previous Thanksgiving	139.6	208.4	-68.8**
<i>Polarization measures:</i>			
Maximum vote share	0.477	0.474	0.003
Standard deviation of vote shares	0.210	0.212	-0.002
Gini index	0.446	0.449	-0.002
Entropy index	-0.725	-0.716	-0.009
<i>Other election variables:</i>			
Years since last election	3.850	3.862	-0.012
Seats per 1000 population	0.101	0.111	-0.010
Number of parties	4.65	4.62	0.029
Voter turnout	0.702	0.684	0.018
Vote share - Conservative	0.325	0.324	0.000
Vote share - Liberal	0.320	0.357	-0.037
Vote share - NDP	0.216	0.177	0.038
<i>Demographics:</i>			
Age, 15 to 24	0.217	0.201	0.017**
Age, 55 plus	0.225	0.241	-0.016
Female	0.510	0.511	-0.002
College degree	0.391	0.412	-0.022
Population in 1000s	6,057.3	2,350.4	3,707.0***
Observations	20	58	78

This table shows the mean of election outcomes, election characteristics, and population characteristics for elections occurring after 1980, but before each jurisdiction adopted fixed-date legislation. Column (1) shows the means for jurisdictions where the eventual fixed date would be just after Thanksgiving, column (2) shows the mean for all other jurisdictions, and column (3) shows the difference between them. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$

Table XI. Effect of fixed-date electoral reform on the timing of elections

	Within 4 weeks of Thanksgiving	Days since Thanksgiving
Treated x Post	0.482** (0.193)	-153.7*** (41.536)
Observations	110	110
R-squared	0.2664	0.1835

This table shows the results from regressions of the timing of elections on “treated” (a variable equal to 1 if a jurisdiction chose a fixed election date in the month following Thanksgiving) interacted with “post” (a variable equal to 1 for elections occurring after electoral reform.) Regressions also include a “post” indicator, jurisdiction fixed effects, a quadratic time trend, controls for demographic characteristics of the population, and day of week fixed effects. Standard errors are clustered at the jurisdiction level. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$

Table XII. The effect of Thanksgiving on agreement in voting

	Max vote share	Standard deviation deviation	Gini index	Entropy index
Treated x post	0.134** (0.0468)	0.0591*** (0.0148)	0.0977*** (0.0273)	0.112*** (0.0301)
Observations	110	110	110	110
R-squared	0.323	0.335	0.202	0.266


This table shows the results from regressions of variables capturing polarization in Canadian elections on “treated” (a variable equal to 1 if a jurisdiction chose a fixed election date in the month following Thanksgiving) interacted with “post” (a variable equal to 1 for elections occurring after electoral reform.) Each dependent variable is a measure of the inequality of vote shares; higher inequality corresponds to *lower* polarization. Regressions also include a “post” indicator, jurisdiction fixed effects, a quadratic time trend, controls for demographic characteristics of the population, and day of week fixed effects. Standard errors are clustered at the jurisdiction level. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$

Table XIII. The effect of Thanksgiving on agreement in voting, separately by treatment jurisdiction

VARIABLES	Max vote share	Standard deviation	Gini index	Entropy index
Saskatchewan x post	0.211*** (0.0322)	0.0930*** (0.0107)	0.146*** (0.0296)	0.186*** (0.0231)
Newfoundland x post	0.181*** (0.0486)	0.0586*** (0.0135)	0.0531 (0.0351)	0.0508* (0.0251)
Federal x post	-0.00296 (0.00858)	0.0203 (0.0116)	0.0879** (0.0313)	0.0915*** (0.0249)
Observations	110	110	110	110
R-squared	0.398	0.367	0.211	0.293

This table shows the results from regressions of variables capturing polarization in Canadian elections on indicators for each of the treatment jurisdictions interacted with “post” (a variable equal to 1 for elections occurring after electoral reform.) Each dependent variable is a measure of the inequality of vote shares; higher inequality corresponds to *lower* polarization. Regressions also include a “post” indicator, jurisdiction fixed effects, a quadratic time trend, controls for demographic characteristics of the population, and day of week fixed effects. Standard errors are clustered at the jurisdiction level. \*\*\* p<0.01, \*\* p<0.05, \* p<0.10

Figure I. The effect of Thanksgiving on discussions with family



This figure shows the coefficients and standard errors from a regression of an indicator for “any discussion with siblings or extended family” on indicators for the number of weeks before/after Thanksgiving and the covariates/fixed effects described in the text. The omitted group is either 3+ weeks before Thanksgiving or 5+ weeks after Thanksgiving. Standard errors are clustered at the survey wave by country level.

Table XIV. Placebo test of effects of Thanksgiving on other election variables

	Seats per 1000 pop.	Years since last election	Number of parties	Voter turnout
Treated x post	0.0198 (0.0260)	-0.0449 (0.412)	-0.254 (0.363)	-0.0813* (0.0430)
Observations	110	110	110	110
R-squared	0.133	0.062	0.288	0.634
	Vote share Conservative	Vote share Liberal	Vote share NDP	
Treated x post	0.00620 (0.131)	-0.0665 (0.0447)	-0.0235 (0.0474)	
Observatoins	110	110	110	
R-squared	0.281	0.299	0.207	

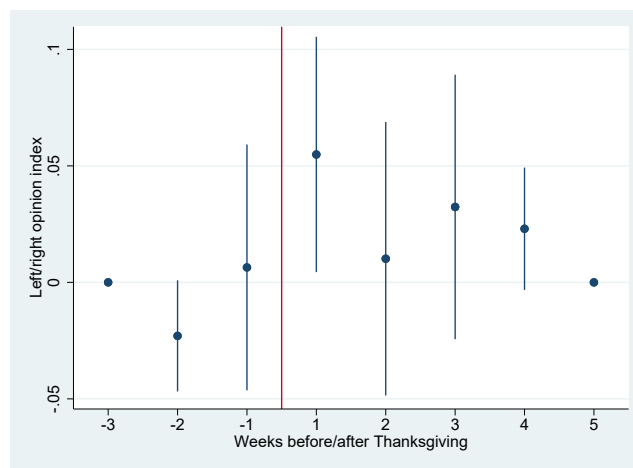
This table shows the results from regressions of other election characteristics and outcomes on “treated” (a variable equal to 1 if a jurisdiction chose a fixed election date in the month following Thanksgiving) interacted with “post” (a variable equal to 1 for elections occurring after electoral reform.) Regressions also include a “post” indicator, jurisdiction fixed effects, a quadratic time trend, and controls for demographic characteristics of the population. Standard errors are clustered at the jurisdiction level. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$

Figure II. The distribution of political beliefs, MTurk respondents and their families



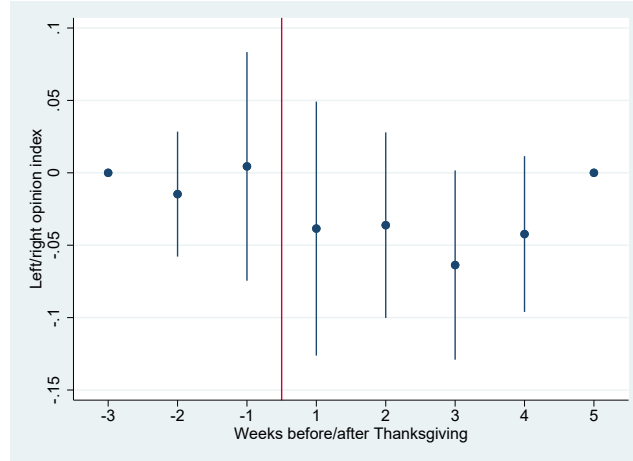
This figure shows a histogram of the reported left/right orientation of MTurk respondents, as well as a kernel density estimate of the distribution of the left/right orientation of their families. The family variable is an average of the reported left/right orientation of siblings and extended family.

Figure III. The effect of Thanksgiving on political opinions (family to the right of respondent)




This figure shows the coefficients and standard errors from a regression of a respondent's left-right political opinion index on indicators for the number of weeks before/after Thanksgiving and the covariates/fixed effects described in the text, for the sample with families who are more right-wing than they are. Family opinions are measured by the average of siblings and extended family. Higher numbers indicate that the individual expresses more right-wing political opinions. The omitted group is either 3+ weeks before Thanksgiving or 5+ weeks after Thanksgiving. Standard errors are clustered at the survey wave by country level.

Figure IV. The effect of Thanksgiving on political opinions (family to the left of respondent)



This figure shows the coefficients and standard errors from a regression of a respondent’s left-right political opinion index on indicators for the number of weeks before/after Thanksgiving and the covariates/fixed effects described in the text, for the sample with families who are more left-wing than they are. Family opinions are measured by the average of siblings and extended family. Higher numbers indicate that the individual expresses more right-wing political opinions. The omitted group is either 3+ weeks before Thanksgiving or 5+ weeks after Thanksgiving. Standard errors are clustered at the survey wave by country level.

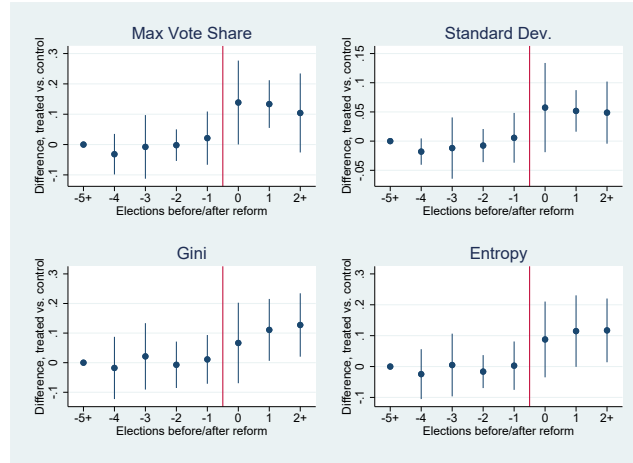
Figure V. The effect of electoral reform on election timing



This figure shows the coefficients and standard errors from a regression of an indicator for whether an election was held within 4 weeks of Thanksgiving on “treated” interacted with indicators for the number of elections before/after a jurisdiction’s reform. The regression also includes the controls described in the text.

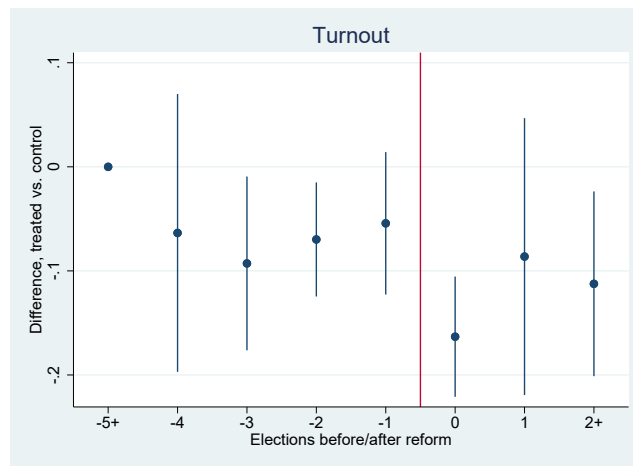


Figure VI. The effect of electoral reform on agreement in elections



This figure shows the coefficients and standard errors from regressions of measures of agreement within the Canadian electorate on “treated” (an indicator for whether a jurisdiction chose a fixed election date in the month after Thanksgiving) interacted with indicators for the number of elections before/after a jurisdiction’s reform. The regression also includes the controls described in the text.

Figure VII. Event study for voter turnout



This figure shows the coefficients and standard errors from a regression of voter turnout in Canadian elections on “treated” (an indicator for whether a jurisdiction chose a fixed election date in the month after Thanksgiving) interacted with indicators for the number of elections before/after a jurisdiction’s reform. The regression also includes the controls described in the text.