

Politics at the Dinner Table: Thanksgiving and Family Influences on Political Opinions

Kirsten Cornelson*

Abstract

Does who we socialize with affect our political opinions? I answer this question using a shock that changes our social interactions: 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 leads to more votes for the most popular political parties in elections. The results suggest that increasing political diversity within our social circles can reduce polarization.

*Department of Economics, University of Notre Dame, 3051 Jenkins Nanovic Hall, Notre Dame, IN 46556 e-mail: 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 (Iyengar et al., 2019). 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).

A popular explanation for these trends in political polarization is an increasing prevalence of “echo chambers”, or exposure to opinions and news sources that reinforce our existing viewpoints (e.g., Sunstein, 2001, 2017). According to the echo chamber hypothesis, the advent of social media, the polarization of traditional media outlets, and the increased geographic sorting of politically similar individuals all increase our exposure to pro-attitudinal influences and contribute to our tendency to polarize (e.g., Parisier, 2011; Jamieson and Capella, 2008; Bishop, 2008). Trends in time use suggest that social interactions in general are less politically diverse than they were in the mid-1960's¹, which would also work to increase polarization according to this viewpoint.

A fundamental assumption of the echo chambers hypothesis is that our political opinions are causally influenced by our social contacts. A large literature in political science and sociology makes this argument (e.g., Berelson et al., 1954; Mutz, 2006; Mason, 2014). Testing this assumption in a causal way is difficult, however, because it requires random or quasi-random changes to individuals' social environments. The most plausible evidence on this point so far comes from random assignment to college dorms, which shows that college roommates tend to be more similar on political preferences after a year of living together (Strother et al., 2021). It is unclear, however, whether or how strongly this result holds in a more general population, or for less intensive forms of interaction.

¹See an early version of this paper, which presented statistics from the American Time Use Surveys showing that we spend significantly more time with spouses and children than we did in 1965 (about 2 hours per day), about the same amount of time with extended family and friends, and significantly less time with neighbors, coworkers outside of work, and any other unrelated adults. Since spouses are very similar to us politically, and coworkers/neighbors are more politically diverse than the rest of our social contacts, these changes imply an decreasing level of exposure to diverse opinions.

Do our social interactions affect our political beliefs and voting behavior? In this paper, I answer this question by using a shock that induces us to socialize and discuss politics with a different group of people than we usually do: Thanksgiving. As I show in my results, Thanksgiving causes us to have more political discussions with parents, siblings, and extended family. Outside of Thanksgiving, we are much more likely to have these discussions with spouses and friends. While this represents a relatively small shock to the political diversity of our interactions on net (our family members are slightly farther away from us politically than our friends and spouses, but still fairly close to us compared to the general population), the exogenous variation in exposure to specific groups of people provides an opportunity to observe what impact these people have on our political opinions. I use two empirical strategies to examine whether we move towards our families' opinions after Thanksgiving, and what consequences this has for political polarization and voting behavior.

The first strategy takes advantage of the fact that Thanksgiving occurs at different times in the U.S. and Canada. I ran repeated weekly surveys of MTurk workers in September-December of 2019 and 2021, in which I asked questions about political conversations and political beliefs. I examine the impact of Thanksgiving on social interactions and opinions by using Americans as a control group around Canadian Thanksgiving and Canadians as a control group around American Thanksgiving. The existence of a plausible control group means that this strategy does not capture unrelated seasonal effects, which would be the case if we examined trends in political opinions in either country individually around Thanksgiving.

The survey shows that, as we might expect, Thanksgiving results in a large change in the probability of discussing politics and current events with family. In contrast to the results of Chen and Rohla (2018), this tendency does not vary significantly with political distance from family in my sample. The effect of Thanksgiving on exposure to people of different political orientations is relatively small on net, although respondents do report slightly higher levels of disagreement with discussion partners in the week after Thanksgiving. Thanksgiving represents a shock in the direction of political information, however, which depends on the political orientation of a respondent's family. People with left-wing families have

significantly more left-wing discussion partners in the week of Thanksgiving, and people with right-wing families have more right-wing discussion partners. If our social contacts affect our political beliefs, we should see these results reflected in political opinions. I show that this is the case. For every point that a respondent's family moves right on the political scale (0-6), the respondent's political opinions become more right wing by about 5% of a standard deviation. The results imply, for example, that someone a very left-wing family and someone with a very right-wing family would diverge by about 0.3 standard deviations. This effect lasts around 2 weeks, with opinions returning to baseline by 3 weeks after Thanksgiving.

Because Thanksgiving has a relatively limited effect on the political diversity of social interactions, it has essentially no net effect on political polarization. While the holiday itself has a minimal impact, the underlying tendency it reveals - that we move towards the viewpoints of our social contacts - suggests that political polarization likely is affected by political homophily. In other words, positive social interactions do have the potential to reduce political polarization, should we find a way to create more of these interactions between people of different political beliefs.

In support of the assumption that Thanksgiving is driving the effects I document, I run placebo regressions using two groups who are less likely to celebrate Thanksgiving: French Canadians, and immigrants. Because Thanksgiving is not traditionally celebrated among French Canadians, or in other countries, Thanksgiving does not represent a change to the social environment for these groups. I show that there is no tendency to move towards family opinions for this sample. I also show that the results are quite similar across both countries and both survey years.

It is important to note that the effects I document could be driven either by talking to family about politics, or simply by socializing with them; I cannot distinguish between these two channels using my data. The fact that changes in political beliefs move in different directions depending on the orientation of family, and are essentially unrelated to a respondent's own political orientation, suggests that the channel must be related to exposure to family and is not a general effect of Thanksgiving itself (driven, for example, by feelings of gratitude or reduced exposure to media.) However, I cannot say from my data

exactly how this exposure effect operates.

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 families influence each others' voting behavior, then we should see more people occupy a jurisdiction's political "center" - the modal political viewpoint - when an election is held immediately after Thanksgiving. To capture this logic, I examine the impact of Thanksgiving on the vote shares of more and less popular parties in each jurisdiction. I look at how the share of votes captured by the first-, second-, and third-or-lower ranked parties changes in the treatment jurisdictions compared to the control jurisdictions. If we see more people voting for more popular parties when elections are held immediately after Thanksgiving, this suggests that the political majority may be influencing marginal voters in its direction. In other words, we would say that Thanksgiving reduces disagreement among the electorate. This is similar conceptually to reduced polarization, although it captures the level of agreement and disagreement in voting rather than opinions.

Figure V shows the results of this exercise. The figure shows that the winning party's vote share increased by around 10 percentage points after the reform compared to the control group. I show that about half of this effect is driven by reduced turnout among supporters of less popular parties, but that there is also a significant increase in the number of votes the winning party receives. The second place party sees a small and insignificant decline in the vote share, while third-place and lower parties receive substantially less support. These results suggest that there is more agreement among the electorate immediately after Thanksgiving, which is consistent with social influences on voting.

The underlying assumption in this exercise is that there is no other reason why voting patterns would change in the treatment jurisdictions, compared to the control jurisdictions, other than the timing of elections. In support of this assumption, I show that polling data taken 1-3 months prior to the election show no differential advantage for the winning party in treatment compared to control jurisdictions; the advantage becomes evident only in late polling data (which is more likely to occur after Thanksgiving in the treatment group). The fact that the results hold in the subset of elections for which polling data is available and widely reported also suggests that the results are not driven by strategic voting. This could occur if, for example, talking to family at Thanksgiving gave voters concrete information about likely election outcomes. The fact that there is already information about the average voter's opinion at this point, however, suggests that this is not likely to be the case. I also show that there are significant changes in voting patterns in all 3 treatment jurisdictions, suggesting that my results are not driven by random events in a single jurisdiction.

Taken together, the survey and election results both suggest that our social contacts have significant effects on our political beliefs and behaviors. While Thanksgiving itself represents a fairly small average shock to both disagreement and polarization in the MTurk sample, the event reveals a strong and significant relationship between the viewpoints of our social contacts and our own political opinions. This suggests that more frequently socializing with people who disagree with us could, in principle, reduce polarization. The voting results suggest that the effect of changing our social interactions on political behavior may be even more pronounced.

Of course, my results capture the effect of social interactions in a particular context. 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 who are different from us politically. A society where relationships like this are more common is likely to be one with less polarization.

2 Related literature

Much of the recent literature on political polarization has focused on the impact of either social media or the role of partisan cable news. This literature contains several well-identified papers showing that these explanations for political polarization have some merit. For example, researchers have run experiments varying the use or content of social media (Allcott et al., 2020; Levy, 2021), showing that this leads to reduced polarization (although see Bail et al. (2018) for a conflicting result), or used plausible instruments affecting exposure to Fox News to show that the network does indeed pull its viewers to the right (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2017). The latter papers are part of a large literature showing that media has a significant influence on our political beliefs (e.g., Gerber et al., 2009; Chiang and Knight, 2011; Enikolopov et al., 2011). Note, however, that the trend of increasing political polarization began in the 1970's, well before the invention of social media or cable news. This suggests that these explanations do not fully capture the causes of increased polarization, although they do seem to play a role in recent years.

With regard to echo chamber effects specifically, some of the most convincing prior evidence comes from the literature on group polarization, a theory arguing that people become more extreme in their opinions when exposed to others who agree with them. This phenomenon has been demonstrated experimentally in the lab several times (e.g., Stoner, 1961; Moscovici et al., 1972; Swol, 2009), although I am unaware of any papers identifying this phenomenon in a field setting. With respect to social influences on political opinions more generally, the closest paper to this one is Strother et al. (2021), which uses

random assignment to college dorms to estimate the impact of roommates on political opinions. They find small but significant effects of being assigned to a more right-wing roommate on reported political ideology at the end of freshman year. Note that their results relate to reported political *orientation*, the student’s self-identified placement on a left-right spectrum, while my results relate to either opinions on specific issues (in the Canada-U.S. comparison) or voting behavior (when looking at the effects of election reform). To the extent that orientation is more closely related to identity, and therefore more stable over time, their results may understate the influence of social interactions on politics. On the other hand, the interaction they study is much more intensive and long-lasting than the one I examine in this paper.

More broadly, this paper relates to a large literature in economics that documents causal effects of social interactions on economic outcomes (e.g., Duflo and Saez, 2003; Bayer et al., 2008, 2009; Carrell et al., 2009; Dahl et al., 2014; Bingley et al., 2021). Perhaps most relevant to the current paper is the literature on “contact theory”, which examines the impact of social homophily on intergroup attitudes and behaviors. This literature suggests that contact with people from different groups does lead to reduced prejudice and more positive attitudes towards members of the outgroup (e.g., Clingingsmith et al., 2009; Paluck et al., 2018; Carrell et al., 2019; Rao, 2019; Lowe, 2021). While my results capture the influence of diverse groups on political opinions, rather than intergroup attitudes, they also speak to the importance of diverse social interactions in easing intergroup tensions.

Finally, with respect to the effect of Thanksgiving itself, Chen and Rohla (2018) use cell phone location data to show that people tend to spend nearly an hour less at Thanksgiving when the home and Thanksgiving dinner location voted for opposite parties during the 2016 election, and that this effect is stronger in places that saw more political advertising. This paper examines the converse relationship to the one I document in this paper: it shows that political differences lead to fewer interactions, while I show that more interactions lead to fewer political differences.

3 Data and empirical strategy

3.1 Canada-U.S. comparison

My first empirical strategy leverages the difference in the timing of Thanksgiving in the U.S. and Canada, using each country as a control group for the other in the week of Thanksgiving. The use of a similar country as a control group around the same time provides assurance that the effects I document are actually due to Thanksgiving and not simply a seasonal effect. If we observe similar changes in social interactions and polarization in Canada in early October and the U.S. in late November, the most plausible explanation is that Thanksgiving is causing these changes.

My exercise will rely on relating changes in the social environment induced by Thanksgiving to the political viewpoints of a respondent's family. Unfortunately, there are no existing publicly available datasets that contain information on family political orientation and social behavior. I therefore rely on survey data, collected through the Amazon Mechanical Turk platform, for this exercise. Previous literature has shown that the MTurk respondent pool tends to produce high-quality data in terms of consistency and test-retest reliability (e.g., Hall et al., 2016; Buhrmester et al., 2011), and while not representative of the U.S. population, is more representative than other common survey pools (Berinsky et al., 2012). Additionally, MTurk workers tend to respond similarly to other survey respondents in experimental interventions (Horton et al., 2011). Of course, there may still be concerns that what is true in the MTurk sample is not true for the population as a whole; this is a key reason why I also use the Canadian election timing strategy.

I begin in this section by describing the data collection process on MTurk and the key outcome variables. I then provide more detail on the empirical strategy and regression equations.

3.1.1 Survey sample

I recruited a set of 1,025 MTurk respondents in the fall of 2019, and a further 661 workers in the fall of 2021.² The respondents first answered a baseline survey, and then answered follow-up surveys weekly for the following 5 weeks. The first set of respondents (Sample 1) were recruited during the week of October 5th in 2019 (9 days before Canadian Thanksgiving). A second set (Sample 2) were recruited during the week of November 17th in 2019 (11 days before American Thanksgiving). I recruited Sample 3 in 2021, during the week of September 28th (14 days before Canadian Thanksgiving), and Sample 4 during the week of November 8th (17 days before American Thanksgiving.) 136 respondents show up in more than one sample. Most respondents within a sample did not answer surveys in every week; the average number of weeks in a sample is 3.9.

To arrive at my final analysis sample, I drop 25 respondents who did not provide information on their own political orientation, and a further 39 workers who did not provide information on the political orientation of any family members. This is because these variables will play a key role in my analysis. I also drop a further 19 respondents who appear in more than one sample but give substantially different reports on the orientation of their family. Finally, in my main regression analysis, I exclude two groups who (as I show) are much less likely to participate in Thanksgiving: French Canadians, and immigrants. I use these observations for a placebo regression exercise later in the paper. There are 34 respondents who are French but not immigrants, 142 respondents who are immigrants but not French, and 8 respondents who are both. After these exclusions, I am left with 1,368 respondents, and 5,770 respondent-week observations.

3.1.2 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)

²I did not run surveys in 2020, because I believed that the ongoing COVID-19 pandemic would lead to significantly less family contact at Thanksgiving.

to “extremely right wing” (coded as 6). Table I shows summary statistics on both demographics and political orientation. The first row shows that the sample is approximately 1/4 Canadian. This is because I was unsuccessful in recruiting the desired number of participants from Canada.³ 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. Respondents are slightly more likely to be white, and less likely to be black or Asian, than the population averages. The sample is also far more educated than the general population, with around 68% holding a bachelor’s or post-graduate degree. The average left/right orientation in the sample is around 2.8, although slightly more left-wing in Canada.

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 II 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 and extended family, with an average difference of around 1.5-1.7 points. Siblings, parents, adult children, and neighbors (under the “other” category) fall somewhere in between.

³MTurk workers outside of the United States are paid in Amazon gift cards rather than in cash, which, in addition to the smaller population, may account for the limited worker pool in Canada.

3.1.3 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 II shows the probability that a respondent reported having a discussion with each group of people in the previous week, averaged across all survey weeks.⁴ About 86% of respondents report having a political discussion with anyone in the past week. The most common conversation partners are friends (44.8%) and spouses/partners (43.3%), followed by parents (35.3%). Siblings and coworkers have lower rates of discussion, at 19-22%, and extended family, adult children, and others, have the lowest rates of discussion at 13.4%, 9.7%, and 2.9%, respectively.

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 7.5. I also ask respondents how much they enjoyed this discussion, where the possible responses are “not at all”, “somewhat”, or “a lot”. 29.7% of respondents report that they enjoyed their most recent discussion a lot, 62.3% said they enjoyed it somewhat, and 8.0% said they did not enjoy it at all.

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 a scale of 1 to 5 (from “strongly disagree” to “strongly agree”). I wrote these statements to reflect typical left- and right-wing positions on a variety of issues, where the issues were chosen for concordance with the party platform positions rated by political scientists at the Comparative Manifesto Project (Krause et al., 2020). An example of a left-wing statement is: “The military budget

⁴These probabilities are not conditional on having the relationship described; for example, uncoupled people are included as zeros in the variable “discussed with spouse”.

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 randomly select a set of statements so that respondents will not see the exact same set of statements every week (which might lead them to give the same responses every week out of habit.)

I construct the weekly left/right opinion 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. In principle, it can range from 0 to 6; in practice, it has a smaller range of 1.6 to 3.7, with a standard deviation of 0.35.

Finally, there are two variables that I added to the survey in 2021 in order to provide more detail on the results and mechanisms.⁵ First, I asked respondents questions measuring their “affective polarization”, or the level of hostility towards members of other political parties (e.g., Iyengar et al., 2019). I measure this in a way that is very similar to that used in the ANES and other political surveys, by asking participants to rate i) their level of warmth towards each of their country’s major political parties on a scale of 0-100, ii) the extent to which they view members of each party as intelligent, also on a scale of 0-100, and iii) the extent to which they view members of each party as altruistic on a scale of 0-100. I then construct variables capturing the differences between each respondent’s ratings of his or her own party (the party they would be most likely to vote for in the next election) and ratings of other parties. In the case of the US, this is simply the difference between the two major parties. In Canada, I construct a weighted average of ratings of other parties, where the weights are proportional to the party’s share in the previous

⁵I thank participants in the Notre Dame brown bag seminar series for these suggestions.

election; this is similar to the procedure followed by Gentzkow (2016). Second, I asked respondents about their news consumption habits. The question text was “In the past week, how frequently did you spend time doing any of the following activities: reading news articles, watching the news on TV, listening to news on the radio, or listening to news on a podcast?” The possible responses were “Every day”, “Most days”, “2-3 times”, “1 day”, or “Never”. I code these into a variable capturing the number of days implied by their response (7, 5, 2.5, 1, or 0.) The mean of this variable is 5.⁶

3.1.4 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:

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

In this equation D_{icw}^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 Thanksgiving_{ct} , which takes the value of 1 if the respondent answers the survey within 7 days after Thanksgiving occurs in their country. W_t are survey week fixed effects and F_i are individual fixed effects (which absorb the effect of being Canadian.) The individual fixed effects ensure that I compare each respondent to his or her own mean in the weeks outside of Thanksgiving. \mathbf{X}_{icw} is a set of individual level controls which include age, race, education and political orientation; these occasionally vary even with individual fixed effects because some individuals appear in more than one survey wave and have changing values of some covariates (e.g. age).⁷

The key independent variable in this regression varies at the country by week level. A natural way to construct the standard errors in this regression would therefore be to cluster at this level. However,

⁶Additionally, I also randomized the order of the questions in 2021. The results presented by year below suggest that this randomization made no difference to the results.

⁷Age, education level, and political orientation vary for some respondents who appear in more than one sample, but usually do so in a way that seems appropriate (e.g. moving up one age bracket or one education level, or shifting one point on the political scale). Race also varies for some participants; because I did not give people the chance to indicate whether they are more than one race, I assume this is due to multi-racial individuals reporting a different race. Excluding these observations makes no difference to the results.

as noted by Cameron and Miller (2015), this procedure can lead to over-rejection of the null hypothesis when the number of clusters is small. Furthermore, the number of treated clusters in my regression is even smaller; of the 48 country-weeks, just 4 are “treated”. The standard solution for a small number of clusters, the wild cluster bootstrap, does not work well in this case (Cameron and Miller, 2015). Instead, I construct the standard errors by adapting a procedure suggested by Conley and Taber (2011). The intuition is to use variation within the untreated clusters to examine the distribution of estimated effect sizes under the null hypothesis. Specifically, I eliminate the week of Thanksgiving from my sample and randomly assign “Thanksgiving” to one of the remaining dates for each country and year. I then run the regression described above. By repeating this procedure, I observe the distribution of the β s when the true effect size is zero, and can calculate the probability that my estimated effect size in the main regression would occur given this distribution.

To examine time trends in the discussion variables, and to ensure that the increase around Thanksgiving is “sharp”, I will show plots based on an adapted version of the above regression:

$$D_{ict}^j = \alpha + \sum_{k \in -2,3} \beta_k T_{ct}^k + \sum_t W_t + \sum_i F_i + \mathbf{X}'_{icw} \gamma + \epsilon_{icw} \quad (2)$$

where T_k is an indicator for being in the k th week, relative to the week just before Thanksgiving. Unfortunately, I typically only have 2 weeks prior to Thanksgiving for the Canadian sample, and 4 weeks after Thanksgiving for the American sample. This limits my ability to examine longer pre-trends, or to examine what happens beyond 4 weeks after Thanksgiving.

In this case, my procedure for constructing the standard errors does not work well; because I am examining “treatment effects” for several weeks simultaneously, I do not have sufficient data to exclude the treatment observations and randomly assign Thanksgiving within the remaining time periods. In this case, I instead use a procedure that produces very similar standard errors to the bootstrap I use in my main regressions: clustering at the country by wave level and constructing standard errors using a jackknife. This procedure slightly *under*-rejects the null hypothesis in the main regressions.⁸

⁸I show the standard errors for my main regressions using this procedure in a robustness exercise appearing in Table XIV.

It is important to note that the discussion regressions should not be interpreted as a “first stage” in the instrumental variables sense, because they do not capture the full effect of seeing family at Thanksgiving. In particular, seeing family may affect us even if we do not discuss politics with them. However, if we do see that people discuss politics with their families, it makes it more plausible that this is part of the mechanism driving any changes in opinions around Thanksgiving. It is important to keep this in mind when discussing the magnitude of the opinion results.

I next examine what impact Thanksgiving has on measures of disagreement with discussion partners. I run regressions exactly analogous to Equation 1, using outcome variables that capture either i) the average absolute difference in the political orientation ratings between a respondent and all discussion partners in that week, ii) the maximum absolute difference in political orientation between a respondent and his or her discussion partners, iii) reported agreement in the most recent discussion, and iv) measures of enjoyment of the most recent discussion.

Next, I examine whether exposure to left- and right-wing opinions changes at Thanksgiving, and how this relates to family political orientation. I calculate the average left/right orientation of all of a respondent’s discussion partners in the previous week and use this as my outcome variable. I relate changes in this variable to family political orientation, which is constructed as a weighted average of the opinions of parents, siblings, and extended family. Weights are based on the change in the probability of discussion with each group around Thanksgiving.⁹ This variable therefore captures the average political orientation of the people to whom the respondent is exposed because of Thanksgiving. I then run the following regression:

$$R_{ict} = \alpha + \beta_1 Thanksgiving_{ct} + \beta_2 Thanksgiving_{ct} * rw_i^{fam} + \delta rw_i^{fam} + \Sigma_t W_t + \Sigma_i F_i + \mathbf{X}'_{icw} \gamma + \epsilon_{icw} \quad (3)$$

⁹Specifically, the weight on group j for person i is $\frac{\beta_{disc}^j * \mathbb{1}_{ij}^{nonmiss}}{\Sigma_j \beta_{disc}^j * \mathbb{1}_{ij}^{nonmiss}}$ where β_{disc}^j is the coefficient from the discussion regressions, and $\mathbb{1}_{ij}^{nonmiss}$ is an indicator that the left-right orientation for group j is non-missing for person i . The latter variable ensures the weights sum to one for each person.

where R_{ict} is a measure of left/right opinions for person i from country c in week t , and rw_i^{fam} is his or her family’s political orientation. The coefficient β_1 captures the change in opinions at Thanksgiving for people who have very left-wing families; the coefficient β_2 shows how this change varies as the family becomes more right-wing. In some versions of this regression, I also allow the effect of Thanksgiving to vary with a respondent’s own political orientation, rw_i^{self} or with the interaction of rw_i^{fam} with rw_i^{self} . The latter version of the regression tells us whether the effect of exposure to family of a particular political orientation differs for respondents who are more or less right-wing. This would occur, for example, if the effect of family depends on whether the respondent is to the left or right of his or her family. I also show the results visually, using a regression equation similar to Equation 2. This tells us whether any change in the week of Thanksgiving persists beyond that week.

One particular issue with respect to the Canadian results is that there were federal elections close to Thanksgiving in both survey years. The 2019 election occurred 1 week following Thanksgiving, while the 2021 election occurred on September 20th, about 3 weeks prior to Thanksgiving. For this reason, I include a time trend measuring the days leading up to and following the federal elections in Canada. I will also show that the results all hold in the U.S., and are quite similar to the results in Canada; combined with the placebo regression showing no response among French Canadians and immigrants (who would be affected by the elections as well), this suggests that the elections are not contributing substantially to my results.

3.2 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 I, the MTurk sample is not representative of the general population. MTurk respondents are younger, more educated, and slightly 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.

In order to address these concerns, I next use a series of Canadian electoral reforms that affected the timing of elections to examine whether Thanksgiving causes people to vote differently. Specifically, if my results from the last section hold in the general population, then we should see more voters move towards each jurisdiction's political "center" when elections occur just after Thanksgiving. This should result in more popular parties capturing more votes. 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.

3.2.1 Background

Canada and its provinces maintain a British-style system of parliamentary democracy, in which members of parliament are elected to represent districts ("ridings"). Each province maintains its own legislature, which is independent of the federal government, and holds elections separately from the rest of the country. Federal and provincial leaders ("Prime Minister" and "Premiers", respectively) are selected based on which party holds the most seats in the corresponding parliament.

Prior to the mid-2000's, election dates at both the federal and provincial level were typically set by the ruling party, under the constraint that they had to hold an election within 5 years of the previous one.¹⁰ This allowed the ruling party to manipulate election timing to increase their odds of victory, and resulted in more elections than were strictly necessary.

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 were to be held on a specified date 4 years after the previous election. Table III shows the years of reform and the election days chosen by each jurisdiction. The days chosen for elections typically follow a rule along the lines of "The 1st Tuesday in November". Since Canadian Thanksgiving also follows this type of rule (it occurs on

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

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. Most of the other jurisdictions initially chose dates that were either just before Thanksgiving, or were in the spring. Four jurisdictions made later amendments to their dates, with two of these jurisdictions (New Brunswick and British Columbia) moving their dates to occur post-Thanksgiving. These two jurisdictions, however, have not yet held an election that complied with their new fixed dates. I therefore treat these provinces as part of the “control group” in the regressions below, but examine the impact of this choice in the robustness section.

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 2011.¹¹ Secondly, most fixed date election legislation contains clauses allowing or requiring elections to be moved if they coincide with a significant religious holiday, a federal election (in the case of provincial legislation), or if the premier resigns. 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 voting against the government on a major bill, which is known as a vote of non-confidence. In my analysis, I focus on elections that are subject to the fixed-date legislation: that is, elections that are not induced by fallen minority governments, and do not fall under situations permitting them or requiring them to be moved. Among this sample, there is compliance with the fixed date 84% of the time.

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

3.2.2 Data

I gather data on election results - vote shares by party - for all provincial and federal elections from 1980-mid 2022 from Wikipedia.com.¹² I also use this source to gather other election information, such as voter turnout and the number of parties, which I will use in regressions examining mechanisms. I start with 124 elections, but eliminate 7 that are fallen minority governments and 6 that were moved legally because of exceptions set out in the fixed date legislation. This results in a final sample of 111 elections.

The first column of Table IV shows summary statistics on the elections sample. Elections typically occur about 3.8 years apart, with election campaigns lasting around 32 days on average. The average number of parties running in a given election is 4.8, with 46.6% of the vote going to the party with the highest vote share on average.¹³ The three major national parties capture about 84% of the vote on average. About 17% of elections in the sample occur in the 4 weeks after Thanksgiving.

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.¹⁴ Means of these variables are also shown in the first column of Table IV.

My empirical strategy will rely on comparing measures of voting behavior 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. Table IV provides evidence that these jurisdictions looked very similar on measures of political agreement, and on other characteristics, prior to the reform. There

¹²I 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.

¹³Because of the “first past the post” system, the party with the highest vote share does not necessarily correspond to the winning party (the one with the highest number of seats.) In practice, however, the two are typically the same.

¹⁴For elections in my sample occurring prior to 1990, I interpolate values using Census information from 1981 and 1986. For elections occurring towards the end of my sample, for which data was not available, I use the last available value of the demographic variables.

is no economically or statistically significant difference in the vote shares of first, second, or lower-ranked parties. There was a pre-existing difference in the timing of elections, with my treatment group already more likely to hold their elections in the month after Thanksgiving. As I will show, however, there was a large increase in this probability after the reforms. There are no major differences between the treatment and control group in other election characteristics. There does seem to be a difference in the age distribution of the treatment and control group, with the population being slightly younger in the treatment group. Additionally, the treatment group has a larger population than the control group, which is because the federal government is included in the treatment group. 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.

Finally, I collect data on opinion polls prior to each election in order to examine the timing of changes in preferences for political parties. Data on opinion polls is available for recent elections in all jurisdictions on Wikipedia.com; this source also contains historical data for the federal government and some large provinces. I supplement this data by searching for reports of opinion polls in historical newspapers on Newspapers.com, and by using data from the Canadian Public Opinion Research Archive. Unfortunately, it is difficult to find historical data on opinion polling for several of the smaller provinces. I am able to find polls taken within 1-3 months prior to the election for 81 elections, and polls taken within the month prior to the election for 83 elections.¹⁵ Within this data, I calculate the share of decided voters who say they will vote for each party, and calculate the maximum share among these parties.

3.2.3 Empirical strategy

My empirical strategy is to compare electoral outcomes across jurisdictions based on whether their fixed date predicts an election after Thanksgiving or at some other point in the year, before and after the fixed date reforms. My main regression equation is:

¹⁵In order to avoid polls taken on the weekend of Thanksgiving, I use the threshold of at least 33 days prior to the election for the early polling results.

$$V_{jt} = \alpha + \beta Treatment_j * Post_t + \sum_k \gamma^k E_k + \sum_j P_j + \mathbf{X}_{jt}' \delta + \epsilon_{jt} \quad (4)$$

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 jurisdictions that initially chose post-Thanksgiving election dates (Newfoundland, Saskatchewan, or the Federal government) and $Post_t$ is an indicator for whether the election occurred after the fixed-date election reform. The variables E_k index the number of elections relative to the fixed-date reform (a more detailed version of the $Post_t$ variable), P_j are jurisdiction fixed effects, and \mathbf{X}_{jt} is a vector of control variables that includes demographics and indicators for every 5-year period in my data.¹⁶

As pointed out by Bertrand et al. (2004), the standard errors in this equation may be biased due to serial correlation in both the outcome variable and the treatment variable. To construct the standard errors, I therefore use a procedure similar to the one I describe for the survey data. Specifically, I eliminate the observations occurring after the reforms, randomly assign each province a placebo “treatment” date within the remaining years, and run the above regression on that data. I repeat this procedure to elicit the distribution of treatment effects under the null hypothesis, and use this to construct my standard errors.

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 related variables such as the number of parties or the length of the election. Finally, I examine whether the treatment group saw a differential change in pre- or early-election polling data, which would have occurred prior to

¹⁶I do not include year fixed effects because most years contain only a single election, which complicates the interpretation of the results. However, I show in the robustness section that the main results continue to hold if I either include year fixed effects or a cubic time trend.

Thanksgiving. I show that this is not the case, suggesting that the voting changes I document emerged towards the end of the election period - which is consistent with the interpretation that Thanksgiving influences voting behavior.

As noted by Goodman-Bacon (2021), difference-in-differences estimates can be biased when there is variation in treatment timing, as in my setting, and when treatment effects vary over time. This bias arises when we compare early adopters to late adopters. To show that this comparison is not driving my results, I also examine what happens when I compare each treatment jurisdiction individually to the control group - the set of provinces that did not adopt post-Thanksgiving legislation. This set of regressions additionally rules out the possibility that a single jurisdiction is driving my results.

4 Results

4.1 Canada-U.S. comparison

4.1.1 Discussions and exposure to different opinions

First, I examine how Thanksgiving affects the probability of discussing politics or current events with different groups. The results are shown in table Table V. The table shows that Thanksgiving has large and significant impacts on the probability of having a discussion with parents, siblings, and extended family. The point estimates for parents and extended family are 7.0 percentage points and 5.6 percentage points, respectively, and are both statistically different from zero at the 5% level. The point estimate for siblings is slightly smaller at 3.6 percentage points, and is statistically significant at the 10% level. There are no significant changes in the probability of talking to any other group. The combined effect on the probability of talking to any family (parents, siblings, or extended family) is 9.1 percentage points, and is statistically significant at the 5% level.

Figure I shows the effects documented in Table V visually, using the summary variable "discussion with any family" as an outcome. Conversations are about 11 percentage points higher in the week of Thanksgiving compared to the week before, although this is due in part to discussions being unusually

low in the prior week; compared to even earlier weeks, the increase is closer to 6 percentage points. As I show later, however, changes in the political content of our discussions are sharper with respect to Thanksgiving.¹⁷

Chen and Rohla (2018) documents that people spend less time at Thanksgiving dinner when they are likely to be politically distant from them, as proxied by geographic voting patterns. If this is the case, it could limit the extent to which Thanksgiving represents a shock to political information. Table VI directly examines whether this pattern is evident in my sample. The first column shows how the summary dependent variable “any discussion with family” changes at Thanksgiving, while columns (2) and (3) show how this changes depending on the absolute political difference between a respondent and his or her family. In column (2), I use the weighted measure of family political preferences described in the empirical strategy section to measure the difference between a respondent and his or her family. Because the weights may be endogenous in this exercise - the probability of discussion with particular groups may be larger when they are, on average, less different from the respondent - I also present results using a simple average of the preferences of family in column (3). In both cases, the coefficient on the interaction term is very close to zero, indicating that respondents who are more different from their families are no less likely to have a political discussion with family at Thanksgiving.

Note that these results are not necessarily inconsistent with the results from Chen and Rohla (2018), as the outcome measures in the two exercises are different. If, for example, people spend less time at Thanksgiving with politically dissimilar family members but have a higher probability of lapsing into a political discussion “per minute”, then this could generate the results I show.

Table VII shows how measures of agreement and disagreement in the respondent’s most recent discussion vary at 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 partner(s) in a given week. This variable shows no significant increase in the week of Thanksgiving. If I instead look at the maximum disagreement with any discussion partner, the coefficient is slightly larger at around 0.1 points, but remains

¹⁷This is because rates of increase in discussions with extended family, who tend to be more different from our usual discussion partners and therefore influence our exposure patterns more, are also sharper relative to Thanksgiving.

statistically insignificant. On the whole, then, our discussion partners in the week of Thanksgiving are not more different from us politically than our partners in any other week. The next columns of the table use measures of reported agreement and enjoyment in the most recent discussion as outcome variables. Again, the coefficients indicate slightly more disagreement, but the effects are relatively small and (except for the outcome “enjoyed the most recent discussion a lot”) insignificant. On net, Thanksgiving seems to have a relatively limited effect on exposure to disagreement.

Although disagreement does not increase substantially around Thanksgiving, the holiday still induces a significant change to our social group. This suggests that the direction of political influence we experience - left-wing or right-wing - may change as well. I next examine whether respondents experience a differential exposure to discussion partners who are left- or right-wing, depending on the political orientation of their families. This might not be the case if, for example, respondents compensate for seeing their families by avoiding friends of the same political orientation. In Table VIII, I show that there is a significant increase in exposure to family political orientation in the week of Thanksgiving. For people whose families are as left-wing as possible, the average orientation of discussion partners in the week of Thanksgiving falls (moves leftward) by about 0.2 points, which is statistically significant at the 5% level. For each point the family moves right, the average change in the orientation of discussion partners moves right by 0.050 points, which is also statistically significant. This implies that someone with a family that is extremely right-wing would see a rightward shift in their discussion partners of about 0.1 points.

In the next three columns of the table, I report the results from regressions that either interact the Thanksgiving indicator with a respondent’s own orientation (in column (3)), with both own and family orientation (in column (4)), or with the interaction of the two (column (5)). I report these mainly for the purposes of interpreting similar regressions in the next table, which looks at movement in political opinions. The results show that there is no general change in political homophily around Thanksgiving - more left-wing respondents do not have significantly more left- or right-wing partners overall during this week. When we simultaneously include interactions of Thanksgiving with family- and own-orientation, the family coefficient becomes larger and more significant, and the own-orientation becomes negative but

still insignificant. The effect of family does not vary significantly by respondents' own political orientation.

Figure II shows how the relationship between left/right wing exposure and family orientation changes in the weeks before, during, and after Thanksgiving. The top figure shows the coefficients on each week, interpreted as the effect for a respondent whose family is as left-wing as possible (family orientation equal to zero), while the bottom figure shows the coefficients on week interacted with family left/right orientation. The left-hand figure shows that there is a sharp decline (leftward movement) in the average orientation of discussion partners for people with very left-wing families, while the right-hand panel shows that this is offset by an equally sharp rightward increase associated with each point that the family moves right.

While the results so far suggest that there is a significant change to our social environment at Thanksgiving, it is possible that the holiday also induces changes in another important source of political information: news consumption. If, for example, we stop watching the news around Thanksgiving, this could cause us to “mean revert” and move in the direction of our families. To assuage this concern, Table IX shows the results from regressions of a respondent's reported weekly news consumption (in days per week) on the Thanksgiving indicator, and the interaction of this indicator with own and family orientation. Note that because this variable was measured in 2021 only, the sample size is smaller than in other regressions. The coefficients in all cases are very close to zero, suggesting that Thanksgiving does not crowd out news consumption.

Of course, another possibility is that the content of our news consumption changes at Thanksgiving. For example, a respondent with very right-wing family may be disproportionately exposed to Fox News while visiting his or her family's home. Unfortunately, I cannot measure this in my data. Note, however, that this is perfectly consistent with my interpretation that family influences our political beliefs - in this case, through the mechanism of influencing our media exposure.

In sum, Thanksgiving represents a shock to our social interactions, causing us to spend more time with our families than usual. For people with left-wing families, this creates an increase in exposure to left-wing opinions, and vice-versa for people with right-wing families. If other people affect our political

beliefs, we should see similar patterns in political opinions. I examine whether this is the case in the next set of results.

4.1.2 Changes in opinions

Table X shows how political opinions change in the week of Thanksgiving, and how this is related to the political orientation of family. The first column shows that there is no net change in the left-right opinion index at Thanksgiving. Column (2) of the table shows how this varies across respondents whose families have different orientations. The results indicate that a respondent whose family is as left-wing as possible moves about 0.047 points, or 13% of a standard deviation, to the left at Thanksgiving. For every point that a family moves right, this changes by 0.017 points. This implies that a person whose family is as right-wing as possible moves 0.055 points, or 16% of a standard deviation, to the right. People whose families are moderate - about 3 on the 6 point scale - show no change in their political opinions at Thanksgiving.

In column (3), I examine whether there is any net effect of Thanksgiving on political polarization by examining how Thanksgiving affects respondents of different political orientations on net. The coefficients indicate no significant change in polarization: while people who are very left-wing move slightly to the left and people who are right-wing move slightly to the right, the coefficients are not statistically significant. This is consistent with the fact that Thanksgiving induces very small changes in the overall level of disagreement in our social interactions.

In column (4), I examine whether allowing the effect of Thanksgiving to vary with a respondent's own orientation affects the main results. The tendency to move towards the opinion of family remains of similar magnitude and significance once we allow the effect of Thanksgiving to vary with the respondent's political orientation. The coefficient on own orientation is very close to zero. This version of the regressions shows that two individuals of the same political orientation move in opposite directions, depending on the orientation of family. This, combined with the lack of any general trends in polarization around Thanksgiving, strongly suggests that the effects are indeed driven by family interactions, and not by any

general effect of Thanksgiving on political beliefs.

Finally, in column (5), I allow the effect of family orientation to vary depending on the individual's own orientation. We might expect, for example, that the effect of meeting a family that is “moderately right-wing” (a 4 on the 0-6 scale) will work in different directions for respondents who are left-wing vs “extremely right-wing”. On the other hand, the group polarization hypothesis would suggest that someone who is “extremely right-wing” may move right, even if they meet someone who is only “right-wing”. By allowing the estimated relationship between family orientation and changes in beliefs to vary with a respondent's own orientation, we can capture these types of relationships in the data. In practice, however, the coefficient on the interaction term is small and statistically insignificant; adding this interaction does not substantially change the regression results.

Figure III shows the results from column (4) of Table X visually. It plots the coefficients on indicators for each week before, during, and after Thanksgiving in the left-hand panel, and the interaction between each week interacted with family right-wing orientation in the right-hand panel. The left-hand panel shows that people with very left-wing families move sharply to the left in the week of Thanksgiving, and remain more left -wing in the following week, before returning to baseline by 2 weeks after Thanksgiving. The right-hand panel shows that the difference in this pattern for people with families who are more right-wing also follows this approximate pattern, although the return to baseline is somewhat faster.

While there seems to be no effect of Thanksgiving on opinion polarization, it is possible that it has an impact on affective polarization - the level of hostility that political partisans feel towards the outgroup (people who support other political parties.) In Table XI, I examine how measures of outgroup antipathy vary at Thanksgiving. The results are consistent with the opinion regressions, showing a small and non-significant increase around Thanksgiving.

4.1.3 Placebo test and robustness

To provide additional evidence that the results I documented in the last section are a causal effect of Thanksgiving, I first examine whether we see similar trends among a group of people that are unlikely

to celebrate Thanksgiving: French Canadians, and immigrants. As shown in column (1) of the top panel of the table, there is no increase (in fact, there is a slight decrease) in the probability of discussing politics or current events with family for these groups. This is consistent with the fact that Thanksgiving is not a traditional holiday in either French Canada or in most immigrants' source countries. The second panel of the table shows that the relationship between family opinions and respondent opinions at Thanksgiving also does not hold in the French/immigrant sample. The coefficients on Thanksgiving and on the interaction of "Thanksgiving" and family orientation are of the wrong sign, although the difference is statistically significant only for the main effect. In other words, among groups that do not see family at Thanksgiving, there is no change in the relationship between family opinions and own opinions at this time.¹⁸

Another potential concern with the results is that they might be driven by events in a single country (e.g., a seasonal effect), a single year (e.g. the ongoing pandemic in 2021), or even a single country-year (e.g., the Canadian 2019 election, which occurred a week after Thanksgiving). In Table XIII, I show the main results separately by country and by year. The results are strikingly similar across samples. Within each country and year, the coefficient on "Thanksgiving" is negative (indicating leftward movement among discussion partners and in opinions for people with very left-wing families), and the coefficient on "Thanksgiving" interacted with family orientation is positive. The effect sizes in the latter case are quite stable across samples, ranging from 0.044-0.080 for the discussion partner regressions and 0.008-0.021 for the opinion regressions. Of course, the standard errors get much larger when breaking the data down in this way, meaning that most of the coefficients are no longer statistically significant. The results suggest, however, that the basic effect holds in each of the four "experiments" contained within my data.

Finally, in Table XIV, I show that the results remain nearly identical when I make alternative data choices or use different specifications. In column (1), I use a simple weighted average of the orientation of parents, siblings, and extended family to measure family orientation, which makes essentially no difference to the results. In column (2), I show that the results are very similar when I omit individual fixed effects.

¹⁸While I do not show this in the table for brevity, the point estimates are extremely similar if I look at either the French or immigrant samples in isolation; neither group appears to see family more often at Thanksgiving, and neither group moves towards their families politically.

This suggests that selection into who answers the survey during the week of Thanksgiving is not a major problem in this context. In column (3), I omit the demographic controls from my regressions. Since these are mostly absorbed by the individual fixed effects, it is unsurprising that the coefficients are very similar to those in the main regressions. In column (4), I omit the time trend leading up to the Canadian election, which again makes very little difference. This provides further evidence that the elections are not contributing substantially to the results. Finally, in columns (5) and (6), I show the standard errors when I use either clustering, or clustering plus a jackknife, instead of bootstrapping the standard errors. While clustering alone leads to standard errors that are slightly too small, the jackknife produces standard errors that are very similar to those in my main regressions.

4.2 Canadian electoral reforms

The previous exercise showed that exposure to family at Thanksgiving affects political opinions 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? In this section, I use the Canadian electoral reform sample to show that Thanksgiving induces meaningful changes in voting behavior, which overcomes both limitations of the previous exercise.

4.2.1 Election timing and voting behavior

Table XV 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 63.4 percentage points in the treatment group, compared to the control group. As shown in Column (2), this corresponds to a decrease of about 136 days since the previous Thanksgiving. Figure IV shows that this increase occurs sharply at the time of the reforms. It also shows that compliance is quite high in the first two elections after the reforms, and declines in third elections and beyond.

Table XVI shows how the vote shares of the first, second, and lower-ranked parties changed in the treatment group, relative to the control group. There is a large and statistically significant 12.2 percentage point increase in the vote share of the party with the most votes. The second-place party vote share decreases by an insignificant 3.4 percentage points, while the share for the lowest-ranked parties declines by 8.8 percentage points. Figure V shows the result for the maximum vote share visually. It shows that there is no differential trend in the maximum vote share prior to the reform, and that there is a large increase immediately after the reform. It also shows that the maximum vote share declines in the third election after the reform, which mirrors the first stage results.

The 12.2 percentage point increase in the maximum vote share is surprisingly large, especially given that the “first stage” effect of the reforms on the timing of elections is not perfect. Note, however, that voting shifts of this size are not completely unprecedented in the literature. For example, 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. I also show next that the results are partially driven by reduced voter turnout, which implies a smaller number of vote “switches” than indicated by the main result.

4.2.2 Mechanisms: voter turnout

The results in Table XVI could be generated by two mechanisms: changes in the number of votes going to the winning party, or declines in turnout among voters who were likely to vote for lower ranked parties. Specifically, the winning party vote share is

$$\frac{Votes^{max}}{Votes} = \frac{\frac{Votes^{max}}{Pop}}{\frac{Votes}{Pop}}$$

where Pop refers to the entire voting-eligible population. If turnout declines disproportionately among supporters of lower-ranked parties, the denominator will fall faster than the numerator, resulting in an increase in the winning party’s vote share.

While the distinction between these two mechanisms does not affect the causal nature of my results, it does potentially affect the interpretation. If the number of votes for the winning party per *eligible* voter

(as opposed to per *actual* voter) increases, this is clear evidence that more voters preferred that party at the time of the election. In contrast, a result driven primarily by reduced turnout is more difficult to interpret. It could be that supporters of lower-ranked parties are less motivated to vote for their party because they agree with them less; alternatively, it could simply be that these voters are discouraged from voting by disagreement with their families, even if they still agree with the party’s platform.

Table XVII shows the results for the maximum vote share, the number of votes for the winning party per capita, and voter turnout separately. The table shows that both mechanisms are at work: while the number of votes for the winning party per capita does show an increase of 3.6 percentage points in the treatment group, there is also a decline in voter turnout of about 5.1 percentage points. The bottom panel of the table estimates the contribution of each effect to my main result by taking logs of each dependent variable; in this form, the changes in the numerator and denominator of the above equation add up to the change in the main outcome variable. This exercise shows that the maximum party’s vote share increased by 22.8% in the treatment group, of which about 60% can be explained by a 13.8% increase in votes per capita for the winning party. The remaining 40% can be explained by a 9.1% decline in voter turnout, which we can infer was stronger among supporters of lower-ranked parties.

4.2.3 Results by area

In Table XVIII, I break down my estimates separately for each of the treatment jurisdictions. This exercise assuages two potential concerns about my results. First, a key limitation of my analysis is the small number of observations. With just 3 jurisdictions in my treatment group, it is possible that random events in one jurisdiction could drive my results. Showing significant changes in all treatment jurisdictions, however, gives us more confidence in the underlying mechanism.¹⁹ Secondly, the exercise also shows that my results are not driven by comparisons between groups that are treated earlier vs later (where “treatment” here means adopting a fixed date after Thanksgiving.) As shown by Goodman-Bacon (2021), these comparisons can be biased when treatment effects change over time. The exercise in

¹⁹Technically, the results could also be driven by an outlier in the control group. This is not the case: my main results come from increases in the treatment group, with no change in pre/post behavior in the control group.

Table XVIII compares each member of the treatment group separately to the control (or “never treated”) group.

The top panel of Table XVIII shows that the change in the probability of holding an election within 4 weeks of Thanksgiving goes up in all 3 jurisdictions, although the effect is much stronger in Newfoundland. However, the more continuous version of the first-stage results, which uses days since the previous Thanksgiving as the outcome variable, tells a slightly different story. The first stage in this case shows large changes for Newfoundland and Saskatchewan (125 days and 190 days, respectively) but a smaller change for the federal government (77 days).

The bottom panel of Table XVIII shows that there are large changes in voting behavior in all three jurisdictions, although the change is smaller and of a slightly different nature for the federal government relative to the two provinces. Newfoundland and Saskatchewan both see 18 percentage point increases in the vote share for the winning party, while there is a small and insignificant 1.4 percentage point change for the federal government. However, federal elections show a 7.2 percentage point increase in the vote share of the second place party, and a 9.1 percentage point decline in the vote share of 3rd place and lower-ranked parties (both marginally significant.) The difference between the federal and provincial results could result from either a somewhat weaker first stage, or from the fact that the shock induced by Thanksgiving is more “diffuse” for the federal government. Because there are strong regional differences in federal voting patterns, and because most Thanksgiving conversations will take place within region, we should expect Thanksgiving to reinforce voting patterns by region. In some areas, this will mean that the party that benefits will be one that takes second or third-place nationally.

4.2.4 Placebo and robustness checks

The key threat to identification in this setting is the possibility that some other factor changed voting patterns in Saskatchewan, Newfoundland, and federal elections at the same time that fixed-date legislation changed the timing of elections in these areas. In this case, the changes I document in the last section might not be due to the effect of holding the election after Thanksgiving. One way to rule out alternative

explanations is to look at whether I see changes in political preferences in these jurisdictions, measured *prior* to Thanksgiving. If we see that these jurisdictions already showed more agreement in political preferences at this time, it suggests that Thanksgiving is not the channel driving my results. Conversely, if these preferences emerge after Thanksgiving, it provides additional support for my interpretation.

Table XIX examines whether the treatment jurisdictions see any increase in support for the winning party prior to Thanksgiving, as measured by opinion polls taken 1-3 months prior to the election. Because the fixed dates for elections tend to be either around Thanksgiving or in the spring, these polls should almost never be taken after Thanksgiving when fixed date laws are observed. The top panel of Table XIX first replicates the main results for voting shares on the reduced sample (81 elections) for which I am able to find polls in the appropriate time frame. It shows that both the first stage and second stage results hold and are statistically significant within this smaller sample. The increase in the probability of holding an election within 4 weeks of Thanksgiving is 44.1 percentage points, and the increase in the winning party's vote share is 9.5 percentage points. These results are not apparent, however, in early polling figures. The early polls are slightly less likely to be held after Thanksgiving (although not significantly so) in the treatment group after the reform, and show a small and insignificant 2.9 percentage point increase in the share of voters who indicate that they will vote for the most popular party.

The second panel of Table XIX repeats this exercise, but uses polls taken within the month leading up to each election. Because the sample of elections for which I have late polling data is slightly different than in the previous exercise, I again replicate the main results in this panel; these are very similar to the early-polling sample. In this case, the polling data shows a significant increase in the probability of being taken after Thanksgiving. The share of polls taken in the 4 weeks after Thanksgiving increases by around 35 percentage points, and the maximum party's share in the polling is 11.5 percentage points higher. This exercise shows that the polling data correlates closely with the voting data, and that the shift in preferences towards the winning party occurs towards the end of the election period.

Note that the results in the first columns of Table XIX also help to mitigate concerns about another mechanism that could explain my results: strategic voting. Suppose, for example, that a supporter of the

NDP (a left-wing party) goes to Thanksgiving dinner and learns that most of his left-wing relatives are planning to vote for the Liberals (a center-left party). This might induce him to switch his vote to the Liberals for strategic reasons, to avoid splitting the left-wing vote. However, when polling numbers are available, the strategic value of information provided by relatives should be minimal. This is particularly true for later elections in my sample, when riding-by-riding polling predictions are typically available; although not shown in the table, the results are slightly stronger if I restrict the data to the last 20 years.

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 XX that there were no other observable changes in the characteristics of elections in the treatment group compared to the control jurisdictions. In particular, there is no change in the number of seats per capita, years between elections, the number of parties involved in an election, or the length of elections.

Finally, in Table XXI, I show that my results are not sensitive to the choices I have made about the sample or the empirical model. In the top left part of the table, I show how my results change when I include British Columbia and New Brunswick as part of the treatment group. As noted in the data section, these jurisdictions initially chose fixed dates in the spring and before Thanksgiving, respectively; however, they both switched their fixed dates to occur post-Thanksgiving in legislation introduced in 2017. Both provinces, however, immediately violated their fixed date legislation and held early elections in 2020. As shown in Table XXI, including these provinces as part of the treatment group reduces both the first and second stage results. The results remain large and statistically significant, however.

In the top right panel of Table XXI, I examine what happens when I include elections that occurred after the fixed date reforms but were moved legally because of provisions in the fixed date legislation (e.g. conflict with a federal election). Because it is unclear in some cases whether the moves were required or simply permitted, it may be more appropriate to consider these as endogenous violations, similar to what happens when a government chooses to ignore fixed date elections. The results show that while the first stage is smaller, both the first and second stage remain large and statistically significant when including

these elections. The second stage is somewhat smaller than in the main sample, but the reduction is smaller in magnitude than the change in the first stage. While this could potentially be a cause for concern, it could also reflect the fact that the legally moved elections typically occur after one or two elections on the fixed date; if there are incumbency effects, then the winning party's electoral advantage may persist even when the fixed date is ignored.

Finally, in the bottom panel of Table XXI, I show that the results are not sensitive to alternative specifications of time effects. The results on the lefthand side show what happens when I include a cubic time trend, and the results on the righthand side show what happens when I include year fixed effects. The results are extremely similar to those in the main sample.

5 Conclusion

The survey and election results both suggest that our social contacts affect our political opinions and voting. While Thanksgiving itself represents a relatively small shock to political homophily and opinion polarization, it nonetheless shows us that who we spend time with matters for these outcomes. This suggests that exposure to people with different political viewpoints - at least, when these interactions occur in a positive social setting - can help reduce polarization, as the echo chambers hypothesis suggests. More broadly, the results suggest that changes in our pattern of social interactions over the past several decades may be contributing to political polarization.

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

6 Tables and Figures

Table I. Summary statistics, MTurk sample

	Both	Canadian	American
Canadian	26.8%	100.0%	0.0%
Age - 18 to 25	14.9%	22.7%	12.0%
Age - 26 to 45	66.2%	65.5%	66.5%
Age - 46 and older	18.9%	11.8%	21.5%
Male	57.7%	56.6%	58.1%
White	81.0%	74.1%	83.5%
Black	7.7%	3.4%	9.3%
Asian	6.5%	14.9%	3.5%
High school or less	10.1%	12.6%	9.2%
Some college	21.4%	19.4%	22.2%
College degree	52.0%	56.8%	50.2%
Post-graduate degree	16.4%	11.2%	18.4%
Left/right orientation (0-6)	2.79	2.42	2.93
Warmth, own vs other party (1-100)*	31.6	38.7	28.9
Intelligence, own vs other party (1-100)*	22.5	25.3	21.5
Altruism, own vs other party (1-100)*	20.1	23.8	18.7
News consumption (days/week)*	5.04	4.97	5.07
Observations	1,368	366	1,002

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).

* 2021 only; sample size is smaller

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

	Discussed in past week	L/R distance
Anyone	85.5%	
Friends	44.8%	1.008
Spouse/partner	43.3%	0.793
Parents	35.3%	1.381
Siblings	22.0%	1.287
Coworkers	19.0%	1.552
Ext. family	13.4%	1.688
Adult children	9.7%	1.038
Other	2.9%	1.387

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 the group “other” refers to people in the respondent’s neighborhood.

Table III. Dates of Canadian electoral reforms

Jurisdiction	Law passed	Fixed date (every 4 years)	First affected election
British Columbia	2001	2nd Tuesday in May	2005
	2017	3rd Saturday in October	2020
Newfoundland	2004	2nd Tuesday in October	2007
Ontario	2005	1st Thursday in October	2007
	2016	1st Thursday in June	2018
Federal	2007	3rd Monday in October	2008
New Brunswick	2007	4th Monday in September	2010
	2017	3rd Monday in October	2020
PEI	2008	1st Monday in October	2011
Saskatchewan	2008	First Monday in November	2011
	2018	Fourth Monday in October	2020
Manitoba	2008	1st Tuesday in October	2011
Alberta	2011	Between March and May	2012
	2021	Last Monday in May	2023
Quebec	2013	1st Monday in October	2018
Nova Scotia	2021	3rd Tuesday in July	2025

This table shows the years in which each jurisdiction in Canada passed legislation affecting the fixed date of its elections, and the day on which the legislation specified that election would occur in the future. Some provinces have multiple rows because they modified the legislation.

Table IV. Summary statistics and balancing test for Canadian elections

	Average (entire sample)	Balance, pre-reform elections		
		Post-Thanksgiving	Group: Other	Difference
<i>Election characteristics:</i>				
Years since last election	3.802	3.850	3.838	0.011
Length of campaign (days)	32.44	31.65	33.16	-1.51
Seats per 1000 population	0.093	0.101	0.107	-0.007
Number of parties	4.784	4.650	4.645	0.005
Voter turnout	0.660	0.702	0.678	0.023
Vote share - Conservative	0.310	0.324	0.320	0.005
Vote share - Liberal	0.331	0.320	0.361	-0.041
Vote share - NDP	0.199	0.216	0.177	0.039
<i>Voting behavior:</i>				
Maximum vote share	0.466	0.476	0.469	0.006
2nd party vote share	0.347	0.342	0.354	-0.011
3rd party and lower vote share	0.187	0.182	0.176	0.006
<i>Election dates:</i>				
Within 4 weeks of Thanksgiving	0.171	0.250	0.081	0.169**
Days since previous Thanksgiving	192.5	139.6	207.9	-68.3**
<i>Demographics:</i>				
Age, 15 to 24	0.190	0.217	0.197	0.020**
Age, 55 plus	0.266	0.225	0.250	-0.025*
Female	0.510	0.510	0.512	-0.002
College degree	0.454	0.391	0.424	-0.034
Population (1000s)	4,159	6,057	2,335	3,722***
Observations	111	20	62	82

This table shows the mean of election outcomes, election characteristics, and population characteristics for Canadian elections. The first column shows the means for all elections, while the second column shows the mean for elections occurring after 1980, but before each jurisdiction adopted fixed-date legislation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$


Table V. The effect of Thanksgiving on political discussions

	Discussion with:				
	Spouse	Parents	Siblings	Extended family	Adult children
Thanksgiving	-0.014 (0.026)	0.070** (0.040)	0.036* (0.025)	0.056** (0.028)	0.010 (0.015)
Observations	5,770	5,770	5,770	5,770	5,770

	Discussion with:			
	Friends	Coworkers	Other	Family - combined
Thanksgiving	0.028 (0.055)	-0.020 (0.017)	0.011 (0.014)	0.091** (0.040)
Observations	5,770	5,770	5,770	5,770

This table shows the results from regressions of variables indicating that the respondent had a political discussion with each group in the past week on an indicator for “Thanksgiving”, individual and survey wave fixed effects, and covariates. Covariates include controls for age, race, education, respondent political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). *** 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 90% confidence intervals from a regression of an indicator for any discussion with family on indicators for each week before/after Thanksgiving, as well as individual and wave fixed effects and covariates. The dependent variable is equal to 1 if a respondent had a discussion with parents, siblings, or extended family. Covariates include controls for age, race, education, respondent political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are constructed using a jackknife estimator with clustering at the wave by country level.

Table VI. Differences in discussion rates by political distance with family

	Discussion with family		
	0.091** (0.040)	0.094** (0.044)	0.086** (0.044)
Thanksgiving			
Thanksgiving x $diff_i$		-0.002 (0.015)	
Thanksgiving x $diff_i^{alt}$			0.003 (0.015)
Observations	5,770	5,770	5,770

This table shows the results from regressions of indicator variables for any discussion with family on an indicator for Thanksgiving, as well as Thanksgiving interacted with the political distance between a respondent and his or her family. The dependent variable is equal to 1 if a respondent had a discussion with parents, siblings, extended family, or adult children. The measure of family political difference in column (2) is based on a weighted average of the political views of these groups, where the weights are proportional to the coefficients from the discussion regressions for each group. In column (3), I take a simple average of the political views of the four groups. Regressions also include the political distance measure alone, individual and wave fixed effects and covariates. Covariates include controls for age, race, education, respondent political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). *** p<0.01, ** p<0.05, * p<0.10

Table VII. The effect of Thanksgiving on discussion characteristics

	Discussion partners		Most recent discussion		
	Mean political distance	Max political distance	Agreement	Enjoyed a lot	Did not enjoy
Thanksgiving	0.014 (0.078)	0.112 (0.093)	-0.196 (0.194)	-0.046* (0.033)	0.016 (0.025)
Observations	4,846	4,846	4,846	4,846	4,846
Mean, dep. var	1.067	1.555	7.320	0.297	0.080

This table shows how measures of agreement or disagreement in respondents' political discussions change in the week of Thanksgiving. They report coefficients from regressions of the indicated dependent variable on an indicator for "Thanksgiving", along with individual and wave fixed effects and covariates. Covariates include controls for age, race, education, respondent political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). The sample includes all individuals who had a discussion in the past week, and who have non-missing values of the dependent variables. *** p<0.01, ** p<0.05, * p<0.10

Table VIII. The effect of Thanksgiving on exposure to left- or right-wing viewpoints

	Dependent variable: average right-wing orientation of discussion partners				
Thanksgiving	-0.029 (0.075)	-0.196** (0.094)	0.018 (0.112)	-0.162** (0.097)	-0.287** (0.167)
Thanksgiving x rw_i^{fam}		0.050** (0.024)		0.073** (0.040)	0.110** (0.051)
Thanksgiving x rw_i^{self}			-0.019 (0.022)	-0.042 (0.033)	0.013 (0.083)
Thanksgiving x rw_i^{fam} x rw_i^{self}					-0.015 (0.018)
Observations	4,846	4,846	4,846	4,846	4,846

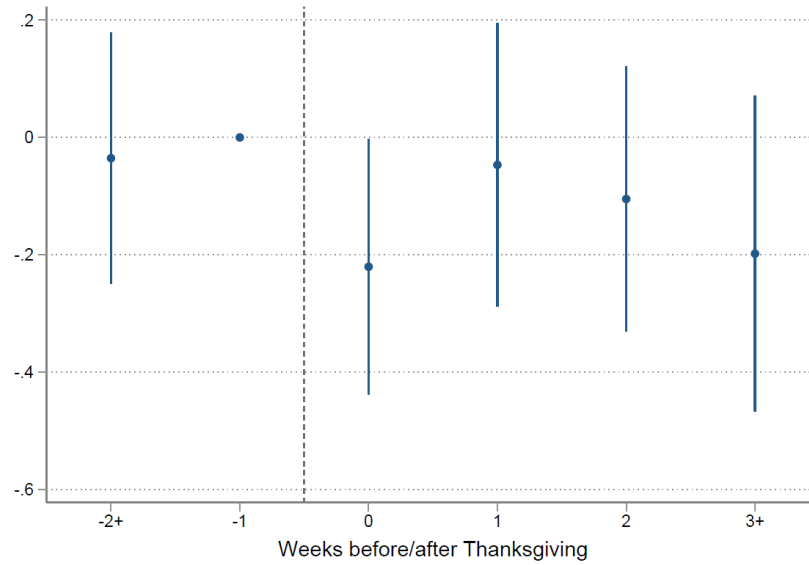
This table shows the results from regressions of the average left/right orientation of a respondent's discussion partners in the past week on indicators for Thanksgiving, and interactions between Thanksgiving and the respondent's family's left/right orientation. In some columns, I also show whether the change at Thanksgiving varies with a respondent's own left/right orientation, or with the interaction of family and own orientation. Regressions also include individual and wave fixed effects. Covariates include controls for the respondent's own political orientation, age, race, education, political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Table IX. The effect of Thanksgiving on news consumption

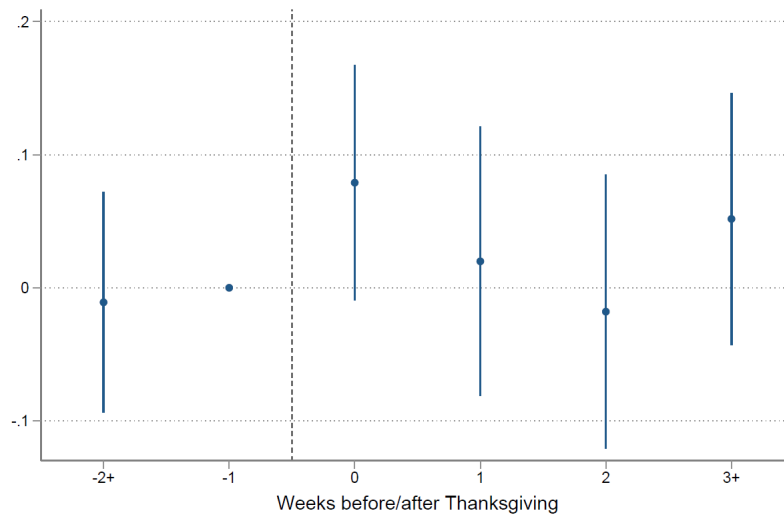
	Dependent variable: news consumption (days per week)				
Thanksgiving	-0.052 (0.093)	-0.038 (0.268)	-0.019 (0.157)	-0.027 (0.297)	-0.038 (0.701)
Thanksgiving x rw_i^{fam}		-0.003 (0.069)		0.004 (0.082)	0.006 (0.241)
Thanksgiving x rw_i^{self}			0.012 (0.040)	-0.012 (0.032)	-0.000 (0.148)
Thanksgiving x rw_i^{fam} x rw_i^{self}					-0.002 (0.047)
Observations	2,285	2,285	2,285	2,285	2,285

This table shows the results from regressions of a respondent's reported news consumption on indicators for Thanksgiving, and interactions between Thanksgiving and the respondent's family's left/right orientation. In some columns, I also show whether the change at Thanksgiving varies with a respondent's own left/right orientation, or with the interaction of family and own orientation. Regressions also include individual and wave fixed effects. Covariates include controls for the respondent's own political orientation, age, race, education, political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$

Figure II. The effect of Thanksgiving on the orientation of discussion partners



(a) Coefficients, $rw_i^{fam} = 0$



(b) Coefficients, week $\times rw_i^{fam}$

This figure shows the coefficients and 90% confidence intervals from a regression of the average left-right orientation of a respondent's discussion partners in the past week on indicators for each week before or after Thanksgiving (in the top subfigure), and these indicators interacted with the family left-right orientation index (bottom subfigure). Regressions include individual and wave fixed effects and covariates, as well as the week indicators interacted with the respondent's own political orientation. Covariates include controls for age, race, education, political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are constructed using a jackknife estimator with clustering at the wave by country level.

Table X. The effect of Thanksgiving on opinions

	Dependent variable: left/right opinion index				
	Both years				
Thanksgiving	0.009 (0.017)	-0.047* (0.033)	-0.005 (0.025)	-0.047* (0.034)	-0.050* (0.037)
Thanksgiving x rw_i^{fam}		0.017** (0.008)		0.017*** (0.007)	0.018* (0.011)
Thanksgiving x rw_i^{self}			0.005 (0.007)	-0.000 (0.007)	0.001 (0.011)
Thanksgiving x rw_i^{fam} x rw_i^{self}					-0.000 (0.003)
Observations	5,770	5,770	5,770	5,770	5,770


This table shows the results from regressions of a respondent's left/right opinion index on indicators for Thanksgiving, and interactions between Thanksgiving and the respondent's family's left/right orientation. In some columns, I also show whether the change at Thanksgiving varies with a respondent's own left/right orientation, or with the interaction of family and own orientation. Regressions also include individual and wave fixed effects. Covariates include controls for the respondent's age, race, education, political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). *** p<0.01, ** p<0.05, * p<0.10

Table XI. The effect of Thanksgiving on affective polarization


	Warmth, own vs other parties	Intelligence, own vs other parties	Altruism, own vs other parties
Thanksgiving	1.047 (1.722)	0.165 (1.731)	1.451 (2.187)
Observations	2,118	2,118	2,118

This table shows the results from regressions of measures of affective polarization on an indicator for Thanksgiving, along with individual and wave fixed effects and covariates. The sample is the set of individuals who answer the survey in 2021, and who have non-missing values of the dependent variable. Covariates include controls for age, race, education, political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are bootstrapped (see text for details). *** p<0.01, ** p<0.05, * p<0.10

Figure III. The effect of Thanksgiving on political opinions



(a) Coefficients, $rw_i^{fam} = 0$



(b) Coefficients, week x $rw_i^{fam} = 0$

This figure shows the coefficients and 90% confidence intervals from a regression of respondents' left-right political opinions index on indicators for each week before or after Thanksgiving (in the top subfigure), and these indicators interacted with the family left-right orientation index (bottom subfigure). Regressions also include individual and wave fixed effects and covariates. The regressions also include the week indicators interacted with the respondent's own political orientation. Covariates include controls for age, race, education, political orientation, and a variable capturing time trends leading up to the Canadian election. Standard errors are constructed using a jackknife estimator with clustering at the wave by country level.

Table XII. Placebo regressions - French and immigrant sample

	Discussion with family		
	Main sample	French/immigrant	Difference
Thanksgiving	0.090*** (0.034)	-0.040 (0.078)	-0.131* (0.081)
Observations	5,770	766	6,536
	Political opinions		
	Main sample	French/immigrant	Difference
Thanksgiving	-0.047* (0.034)	0.095 (0.080)	0.142** (0.085)
Thanksgiving x rw_i^{fam}	0.017*** (0.007)	-0.009 (0.024)	-0.026 (0.026)
Observations	5,770	766	6,536

This table replicates the main results on political discussions and political opinions, using the sample of respondents who are either French Canadian or immigrants. The first column repeats results from Table V and column (4) of Table X, while the second column shows the results of these regressions in the placebo sample. The third column shows the difference between the two samples. All other regression features are the same as described in the corresponding tables for the main sample. Standard errors are bootstrapped (see text for details). *** p<0.01, ** p<0.05, * p<0.10

Table XIII. Results by country and year

	Canada					
	2019		2021		Both years	
	Orientation, discussion partners	Left/right opinions	Orientation, discussion partners	Left/right opinions	Orientation discussion partners	Left/right opinions
Thanksgiving	-0.317 (0.361)	-0.008 (0.130)	-0.204 (0.624)	-0.065 (0.083)	-0.351 (0.290)	-0.011 (0.085)
Thanksgiving x rw_i^{fam}	0.044 (0.074)	0.015 (0.019)	0.078 (0.174)	0.019 (0.020)	0.058 (0.068)	0.014 (0.015)
Observations	880	1,021	423	616	1,303	1,637
	U.S.					
	2019		2021		Both years	
	Orientation, discussion partners	Left/right opinions	Orientation, discussion partners	Left/right opinions	Orientation discussion partners	Left/right opinions
Thanksgiving	-0.048 (0.258)	-0.059 (0.067)	-0.219 (1.610)	-0.148 (0.163)	-0.076 (0.355)	-0.068 (0.085)
Thanksgiving x rw_i^{fam}	0.080** (0.042)	0.021* (0.013)	0.048 (0.581)	0.008 (0.061)	0.078** (0.036)	0.016* (0.010)
Observations	2,089	2,464	1,475	1,669	3,564	4,113
	Both countries					
	2019		2021		Both years	
	Orientation, discussion partners	Left/right opinions	Orientation, discussion partners	Left/right opinions	Orientation discussion partners	Left/right opinions
Thanksgiving	-0.116 (0.147)	-0.026 (0.052)	-0.228 (0.158)	-0.069* (0.052)	-0.141* (0.104)	-0.037 (0.035)
Thanksgiving x rw_i^{fam}	0.068** (0.038)	0.017* (0.011)	0.061 (0.077)	0.013 (0.015)	0.073*** (0.029)	0.015** (0.008)
Observations	2,969	3,485	1,898	2,285	4,867	5,770

This table shows the results from regressions of average discussion partner orientation and the weekly left/right opinion index on an indicator for Thanksgiving, and this indicator interacted with family right-wing orientation, separately by country and year. Family right-wing orientation is measured using year and country-specific weights (note that this is slightly different from the measure used in other tables). The regression also includes an interaction of the Thanksgiving indicator with own right-wing orientation, as well as all of the fixed effects and controls described in the corresponding tables for the main sample. Standard errors are bootstrapped (see text for details). *** p<0.01, ** p<0.05, * p<0.10

Table XIV. The effect of Thanksgiving on opinions: robustness to alternative specifications

	Dependent variable: left/right opinion index					
Thanksgiving	-0.044*	-0.061	-0.048*	-0.047*	-0.047**	-0.047
	(0.033)	(0.061)	(0.034)	(0.033)	(0.019)	(0.032)
Thanksgiving x rw_i^{fam}	0.017***	0.023**	0.018***	0.017***	0.017***	0.017***
	(0.007)	(0.011)	(0.007)	(0.007)	(0.004)	(0.006)
Observations	5,770	5,770	5,770	5,770	5,770	5,770
Specification	Alt. family orientation measure	No indiv. FE	No dem. ctrls	Drop election time trend	Clustered SEs	Clustered jackknife SEs

This table shows how the main results on opinions change when I use slightly different specifications. The first column uses a simple average of parents, siblings, and extended family to measure family political orientation. The second column omits individual fixed effects from the regression. The third column omits the demographic controls, and the third column drops the variable capturing time trends leading up to and following the Canadian elections. In columns (5) and (6), I show the standard errors from either i) clustering at the country by survey wave level, or ii) clustering at the country by survey wave level and using a jackknife estimator. All other regression details are the same as in Table X.

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

	Within 4 weeks of Thanksgiving	Days since Thanksgiving
Fixed date - post-Thanksgiving	0.634*** (0.270)	-136.4** (71.6)
Observations	111	111

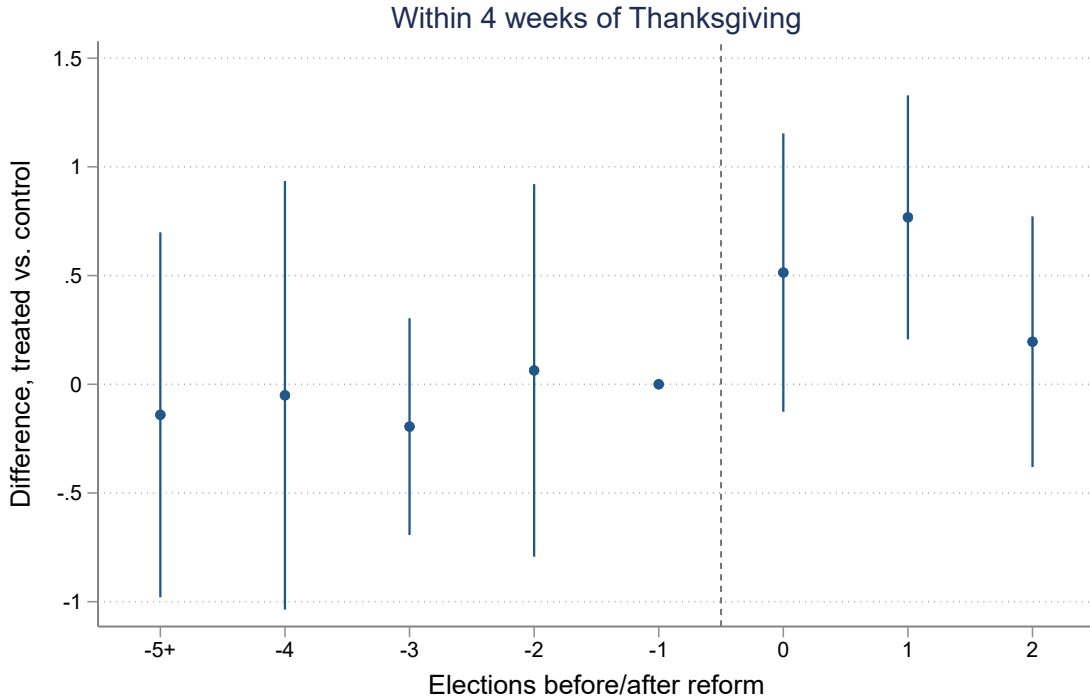
This table shows the effect of fixed-date legislation on the timing of elections in Canada. It shows the results from a regression of the indicated dependent variable on a variable indicating whether fixed date legislation requiring an election within the 4 weeks after Thanksgiving is in place, as well as jurisdiction fixed effects and indicators for the number of elections pre- or post-fixed date reform for all jurisdictions. The regression also includes controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction, as well as indicators for each 5-year period. Standard errors are bootstrapped (see text for details.) *** p<0.01, ** p<0.05, * p<0.10

Table XVI. The effect of Thanksgiving on voting

	Voting shares		
	First place	Second place	Third+ place
Fixed date - post-Thanksgiving	0.122*** (0.030)	-0.034 (0.037)	-0.088*** (0.034)
Observations	111	111	111

This table shows the results from a regression of the vote share of the first-place, second-place, or lower-ranked parties on a variable indicating whether fixed date legislation requiring an election within the 4 weeks after Thanksgiving is in place, as well as jurisdiction fixed effects and indicators for the number of elections pre- or post-fixed date reform for all jurisdictions. The regression also includes controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction, as well as indicators for each 5-year period. Standard errors are bootstrapped (see text for details.) *** p<0.01, ** p<0.05, * p<0.10

Figure IV. The effect of electoral reform on election timing



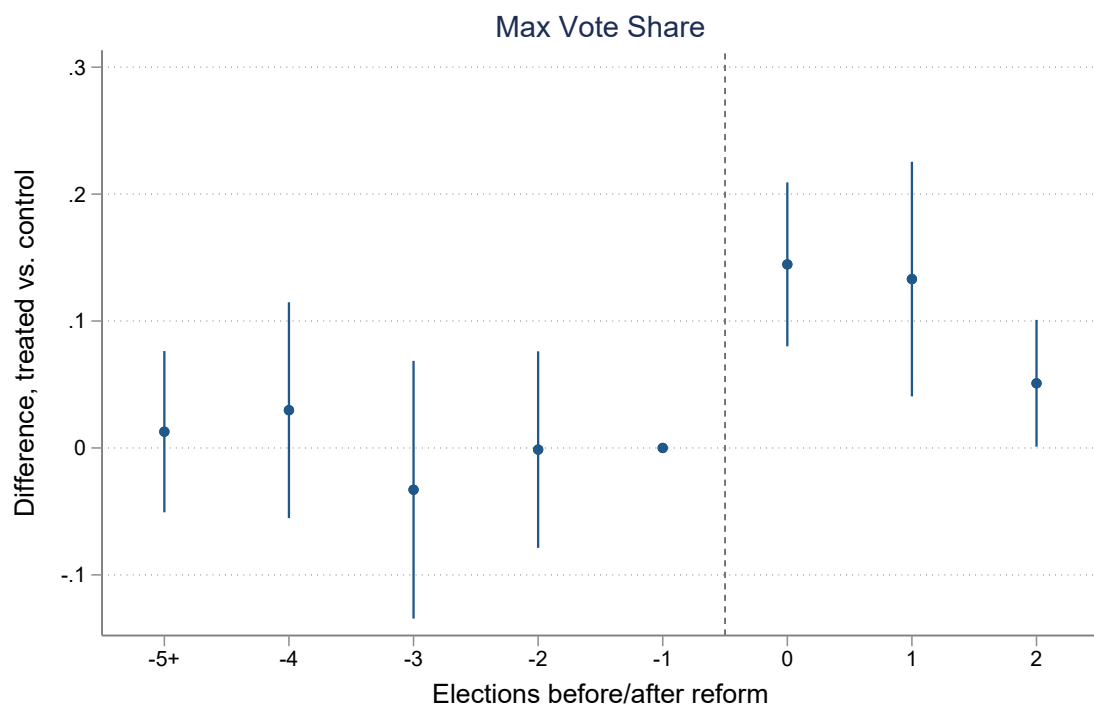
This figure shows the coefficients and 90% confidence interval from a regression of an indicator for whether an election was held within 4 weeks of Thanksgiving on “treated” interacted with indicators for the number of elections before/after a jurisdiction’s reform. The regression also includes jurisdiction fixed effects, indicators for the number of elections pre- or post-fixed date reform for all jurisdictions controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction, as well as indicators for each 5-year period. Standard errors are constructed using a jackknife estimator with clustering at the wave by country level.

Table XVII. Mechanisms - voter participation

	Variables in levels		
	Max vote share	Max votes per capita	Turnout
Fixed date, post-Thanksgiving	0.122*** (0.028)	0.036* (0.026)	-0.051 * (0.033)
	Variables in logs		
	Max vote share	Max votes per capita	Turnout
Fixed date, post-Thanksgiving	0.228*** (0.061)	0.138** (0.084)	-0.091** (0.050)
Observations	111	111	111

This table shows the results from a regression of the indicated dependent variables capturing the first-place party’s vote share and voter turnout on a variable indicating whether fixed date legislation requiring an election within the 4 weeks after Thanksgiving is in place, as well as jurisdiction fixed effects and indicators for the number of elections pre- or post-fixed date reform for all jurisdictions. The regression also includes controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction, as well as indicators for each 5-year period. Standard errors are bootstrapped (see text for details.) *** p<0.01, ** p<0.05, * p<0.10

Figure V. The effect of electoral reform on agreement in elections (maximum party vote share)



This figure shows the coefficients and 90% confidence interval from a regression of the first-place party’s vote share on “treated” interacted with indicators for the number of elections before/after a jurisdiction’s reform. The regression also includes jurisdiction fixed effects, indicators for the number of elections pre- or post-fixed date reform for all jurisdictions controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction, as well as indicators for each 5-year period. Standard errors are constructed using a jackknife estimator with clustering at the wave by country level.

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

	Election timing	
	Within 4 weeks of Thanksgiving	Days since Thanksgiving
Federal x post	0.555* (0.417)	-77.0 (160.0)
SK x post	0.457 (0.429)	-152.9 (127.0)
NL x post	0.966*** (0.344)	-190.3** (87.7)

	Voting shares		
	First place	Second place	Third+ place
Federal x post	0.014 (0.026)	0.072* (0.052)	-0.086* (0.057)
SK x post	0.182*** (0.038)	-0.033 (0.046)	-0.149*** (0.052)
NL x post	0.181*** (0.051)	-0.192** (0.087)	0.010 (0.063)
Observations	111	111	111

This table shows the results from a regression of election timing and vote share variables on interactions between indicators for each treatment jurisdiction and an indicator for the post-reform period. The regressions also include jurisdiction fixed effects, indicators for the number of elections pre- or post-fixed date reform for all jurisdictions, along with controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction and indicators for each 5-year period. Standard errors are bootstrapped (see text for details.) *** p<0.01, ** p<0.05, * p<0.10

Table XIX. Placebo test: early vs late opinion polls

	Early polls			
	Election within 4 weeks of Thanksgiving	Max vote share	Poll within 4 weeks of Thanksgiving	Max poll share
Fixed date - post-Thanksgiving	0.441** (0.258)	0.095*** (0.027)	-0.071 (0.125)	0.029 (0.046)
Observations	81	81	81	81

	Late polls			
	Election within 4 weeks of Thanksgiving	Max vote share	Poll within 4 weeks of Thanksgiving	Max poll share
Fixed date - post-Thanksgiving	0.402** (0.199)	0.102*** (0.038)	0.356** (0.174)	0.115** (0.058)
Observations	84	84	84	84

This table shows the results from a regressions of either voting data or polling data on a variable indicating whether fixed date legislation requiring an election within the 4 weeks after Thanksgiving is in place. The top panel uses polling data from 1-3 months prior to the election, while the bottom panel using polling data between 0-28 days before the election. The regressions also include jurisdiction fixed effects, indicators for the number of elections pre- or post-fixed date reform for all jurisdictions, along with controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction and indicators for each 5-year period. Standard errors are bootstrapped (see text for details.) *** p<0.01, ** p<0.05, * p<0.10

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

	Seats per 1000 pop.	Years since last election	Number of parties	Election length (days)
Fixed date - post-Thanksgiving	-0.026 (0.051)	-0.045 (0.383)	-0.325 (0.441)	4.144 (3.460)
Observations	111	111	111	111

This table shows the results from a regression of the indicated dependent variables capturing other election characteristics on a variable indicating whether fixed date legislation requiring an election within the 4 weeks after Thanksgiving is in place, as well as jurisdiction fixed effects and indicators for the number of elections pre- or post-fixed date reform for all jurisdictions. The regression also includes controls for the population, age distribution, sex distribution, and educational attainment of the jurisdiction, as well as indicators for each 5-year period. Standard errors are bootstrapped (see text for details.) *** p<0.01, ** p<0.05, * p<0.10

Table XXI. Robustness to alternative data or modelling choices

	Add NB/BC to treatment group		Add legally moved elections	
	Within 4 weeks	Max. vote share	Within 4 weeks	Max. vote share
Fixed date - post-Thanksgiving	0.451*** (0.164)	0.094*** (0.025)	0.364** (0.168)	0.115*** (0.027)
Observations	111	111	117	117
	Cubic time trend		Year fixed effects	
	Within 4 weeks	Max. vote share	Within 4 weeks	Max. vote share
Fixed date - post-Thanksgiving	0.571** (0.239)	0.139*** (0.033)	0.557* (0.354)	0.105*** (0.033)
Observations	111	111	111	111

This table shows how the results of the election timing and vote share regressions change when I make alternative sampling or modelling decisions. All other regression details are the same as in Table XVI. *** p<0.01, ** p<0.05, * p<0.10

7 Appendix

7.1 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.
- Allcott, H., L. Braghieri, S. Eichmeyer, and M. Gentzkow (2020). The welfare effects of social media. *American Economic Review* 110(3), 629–676.
- Bail, C. A., L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. F. Hunzaker, J. Lee, M. Mann, F. Merhout, and A. Volfovsky (2018). Exposure to opposing views on social media can increase political polarization. *PNAS* 115(37), 9216–9221.
- Bayer, P., R. Hjalmarsson, and D. Posen (2009). Building criminal capital behind bars: peer effects in juvenile corrections. *Quarterly Journal of Economics* 124(1), 105–147.
- Bayer, P., S. L. Ross, and G. Topa (2008). Place of work and place of residence: informal hiring networks and labor market outcomes. *Journal of Political Economy* 116(6), 1150–1196.
- Berelson, B. R., P. F. Lazarsfeld, and W. N. McPhee (1954). *Voting: a study of opinion formation in a presidential campaign*. University of Chicago Press.
- Berinsky, A. J., G. A. Huber, and G. S. Lenz (2012). Evaluating online labor markets for experimental research: Amazon.com’s Mechanical Turk. *Political Analysis* 20, 351–368.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust difference-in-differences estimates? *Quarterly Journal of Economics* 119(1), 249–275.
- Bingley, P., P. Lundborg, and S. V. Lyk-Jensen (2021). Brothers in arms: spillovers from a draft lottery. *Journal of Human Resources* 56, 225–268.
- Bishop, B. (2008). *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart*. Houghton Mifflin.

- Boxell, L., M. Gentzkow, and J. M. Shapiro (2020). Cross-country trends in affective polarization. *Working paper*.
- Buhrmester, M., T. Kwang, and S. D. Gosling (2011). Amazon’s Mechanical Turk: a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science* 6(1), 3–5.
- Cameron, A. C. and D. L. Miller (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources* 2, 317–372.
- Carrell, S. E., R. L. Fullerton, and J. E. West (2009). Does your cohort matter? Measuring peer effects in college achievement. *Journal of Labor Economics* 27(3), 439–464.
- Carrell, S. E., M. Hoekstra, and J. E. West (2019). The impact of college diversity on behavior toward minorities. *American Economic Journal: Economic Policy* 11(4), 159–182.
- Chen, M. K. and R. Rohla (2018). The effect of partisanship and political advertising on close family ties. *Science* 360, 1020–1024.
- Chiang, C. and B. Knight (2011). Media bias and influence: evidence from newspaper endorsements. *Review of Economic Studies* 78(3), 795–820.
- Clingingsmith, D., A. I. Khwaja, and M. Kremer (2009). Estimating the impact of the Hajj: religion and tolerance in Islam’s global gathering. *Quarterly Journal of Economics* 124(3), 1133–1170.
- Conley, T. G. and C. R. Taber (2011). Inference with “difference-in-differences” with a small number of policy changes. *The Review of Economics and Statistics* 93(1), 113–125.
- Dahl, G. B., A. R. Kostol, and M. Mogstad (2014). Family welfare cultures. *Quarterly Journal of Economics* 129(4), 1711–1752.
- DellaVigna, S. and E. Kaplan (2007). The Fox News effect: media bias and voting. *Quarterly Journal of Economics* 122(3), 1187–1234.

- Dufo, E. and E. Saez (2003). The role of information and social interactions in retirement plan decisions: evidence from a randomized experiment. *Quarterly Journal of Economics* 118(3), 815–842.
- Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and political persuasion: evidence from Russia. *American Economic Review* 101(7), 3253–3285.
- Gentzkow, M. (2016). Polarization in 2016. *White paper*.
- Gerber, A. S., D. Karlan, and D. Bergen (2009). Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics* 1(2), 35–52.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2), 254–277.
- Hall, K. A., D. Cek, A. M. Shaw, and K. R. Timpano (2016). The importance of assessing clinic phenomena in Mechanical Turk research. *Psychological Assessment* 28, 684–691.
- Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011). The online laboratory: conducting experiments in a real labor market. *Experimental Economics* 14, 399–425.
- Iyengar, S., Y. Lelkes, M. Levendusky, N. Malhotra, and S. J. Westwood (2019). The origins and consequences of affective polarization in the United States. *Annual Review of Political Science* 22, 129–146.
- Jamieson, K. H. and J. N. Capella (2008). *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment*. Oxford University Press.
- Kevoian, 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.
- Krause, W., P. Lehmann, J. Lewandowski, T. Matthieb, N. Merz, S. Regal, and A. Werner (2020). Comparative Manifesto Project, version 2020-5.

- Levy, R. (2021). Social media, news consumption, and polarization: evidence from a field experiment. *American Economic Review* 111(3), 831–70.
- Lowe, M. (2021). Types of contact: a field experiment on collaborative and adversarial caste integration. *American Economic Review* 111(6), 1807–1844.
- Martin, G. J. and A. Yurukoglu (2017). Bias in cable news: persuasion and polarization. *American Economic Review* 107(9), 2565–2599.
- Mason, L. (2014). I respectfully disagree: the differential effects of partisan sorting on social and issue polarization. *American Journal of Political Science* 59(1), 128–145.
- Moscovici, S., W. Doise, and R. Dulong (1972). Studies in group decisions ii: Differences of positions, differences of opinion, and group polarization. *European Journal of Social Psychology* 2(4), 385–399.
- Mutz, D. C. (2006). *Hearing the other side: deliberative vs participatory democracy*. Cambridge University Press.
- Paluck, E. L., S. A. Green, and D. P. Green (2018). The contact hypothesis re-evaluated. *Behavioural Public Policy* 3(2), 129–158.
- Pariser, E. (2011). *The Filter Bubble: What the Internet is Hiding From You*. Penguin.
- Rao, G. (2019). Familiarity does not breed contempt: generosity, discrimination, and diversity in Delhi schools. *American Economic Review* 109(3), 774–809.
- Stoner, J. (1961). A comparison of individual and group decisions involving risk. *Massachusetts Institute of Technology Master's Thesis*.
- Strother, L., S. Piston, E. Golberstein, S. E. Gollust, and D. Eisenberg (2021). College roommates have a modest but significant influence on each other's political ideology. *Proceedings of the National Academy of Sciences* 118(2), 1–5.
- Sunstein, C. R. (2001). *Republic.com*. Princeton University Press.

Sunstein, C. R. (2017). *Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.

Swol, L. M. V. (2009). Extreme members and group polarization. *Social Influence* 4(3), 185–199.