

APPENDICES

Contents

A	Case context	45
A.1	Timeline	45
A.2	Correlates of beliefs in systemic racism	47
A.3	Public mourning after the discoveries	48
B	Quasi-experiment	50
B.1	Balance checks	50
B.2	Complete covariate estimates	53
B.3	Belief in systemic racism by survey date	54
B.4	Effects on individual survey items	55
B.5	Effects on racial minorities feeling thermometer	56
B.6	Placebo test: Attitudes toward immigrants	57
C	Over-time analyses	57
C.1	Sample details	57
C.2	Changes in sample composition over time	58
C.3	Systemic racism items separately over time	61
C.4	Response-level frequency over time	62
C.5	Indigenous feeling thermometer over time	63
C.6	Racial minorities feeling thermometer over time	64
C.7	Public prioritization of Indigenous issues over time	65
D	Media analysis	68
D.1	Data sources and text pre-processing	68
D.2	Topic model estimation and validation	69
D.3	Residential school topic prevalence during quasi-experiment	72
D.4	Residential schools and unmarked graves media coverage, 2000 to 2022	72
D.5	Framing analysis	74

E	Canada’s First National Day for Truth and Reconciliation	81
E.1	Balance checks	81
E.2	Media coverage around NDTR	83
E.3	Effects of NDTR on beliefs in systemic racism	84
F	Heterogeneous responses	85
F.1	Partisan effect heterogeneity in quasi-experiment	85
F.2	Quasi-experiment causal forest	86
F.3	Local demographic CATEs in quasi-experiment	90
F.4	Heterogeneity in attitudinal persistence	91

A. CASE CONTEXT

A.1 Timeline

Table A1: Timeline of events, 2021 to 2022

Date(s)	Event
May 27, 2021	215 unmarked graves discovered at former Kamloops Indian Residential School
May 30, 2021	Prime Minister Trudeau orders Canadian Flags to be flown at half-mast
June 21 to July 9, 2021	Over 15 Christian churches are targets of arson attacks by unknown vandals
June 24 to July 8, 2021	Unmarked grave discoveries announced at schools in Marieval (751), Cranbrook (182) and Kuper Island (160)
July 1, 2021	Canada Day celebrations cancelled or scaled back in several regions; counter-celebrations organized by Indigenous activists and allies
August 15 to September 20, 2021	The 44th Canadian federal election campaign results in the re-election of the incumbent Liberal government
September 30, 2021	Canada's first National Day for Truth and Reconciliation is celebrated
January 25 to June 6, 2022	Additional unmarked grave sites discovered at seven former residential schools
March 28 to April 1, 2022	Delegation of Indigenous leaders travel to the Vatican to request an official papal apology
July 24 to 30, 2022	Pope Francis visits Canada and apologizes for the Catholic Church's role in the residential school system

Table A2: List of unmarked grave announcements since 2021

School	Date	Potential graves	Notes
Kamloops, BC	May 27, 2021	200	Initial announcement of 215 potential graves later revised to 200.
Brandon, MB	June 20, 2021	104	Searches were conducted in 2018 and 2019, but discoveries weren't widely covered until June 2021.
Marieval, SK	June 24, 2021	751	
Cranbrook, BC	June 30, 2021	182	
Kuper Island, BC	July 12, 2021	160+	
Williams Lake, BC	January 12, 2022	93	
Fort Pelly, SK	February 14, 2022	42	
St. Philip's, SK	February 14, 2022	12	
Grouard, AB	March 1, 2022	169	
Gordon's, SK	April 20, 2022	14	
Blue Quill's, AB	May 17, 2022	Unknown	Local band announced accidental discoveries of human remains believed to be unmarked graves of former residential school students.
Sandy Bay, MB	May 29, 2022	13	No announcement has been made, but the number of potential graves found is listed in media covering the ongoing searches at the school.
Fort Alexander, MB	June 6, 2022	190	
Pine Creek, MB	June 9, 2022	6	
Pine Creek, MB	August 8, 2022	14	Separate search from the line above.
Lebret, SK	January 12, 2023	2,000	Search also uncovered a child's jaw bone fragment estimated to be 125 years old.
St. Mary's, ON	January 17, 2023	171	
Williams Lake, BC	January 25, 2023	66	

A.2 Correlates of beliefs in systemic racism

In Figure A1, I summarize the partial correlation between a variety of covariates and belief in systemic racism before the first grave discovery was announced (see figure notes for model details). Overall, the model reveals several important patterns. First, partisanship is an important correlate of these beliefs: Conservatives report more than a half standard-deviation lower belief in systemic racism compared to Liberals, and even more compared to other left-wing parties. The size of this difference is comparable to the gap between Indigenous and non-Indigenous people on this issue.

The model also highlights that those who are older, earn more, are men, are distrustful of the media, and live in the West or Atlantic provinces or areas with a larger Indigenous population tend to believe less in the existence of systemic racism. Finally, while Catholics report lower belief in systemic racism than non-religious people, their attitudes are not significantly different from other Christians or followers of other religions. That being said, all of the differences described in this paragraph pale in comparison to the huge partisan gap on this question.

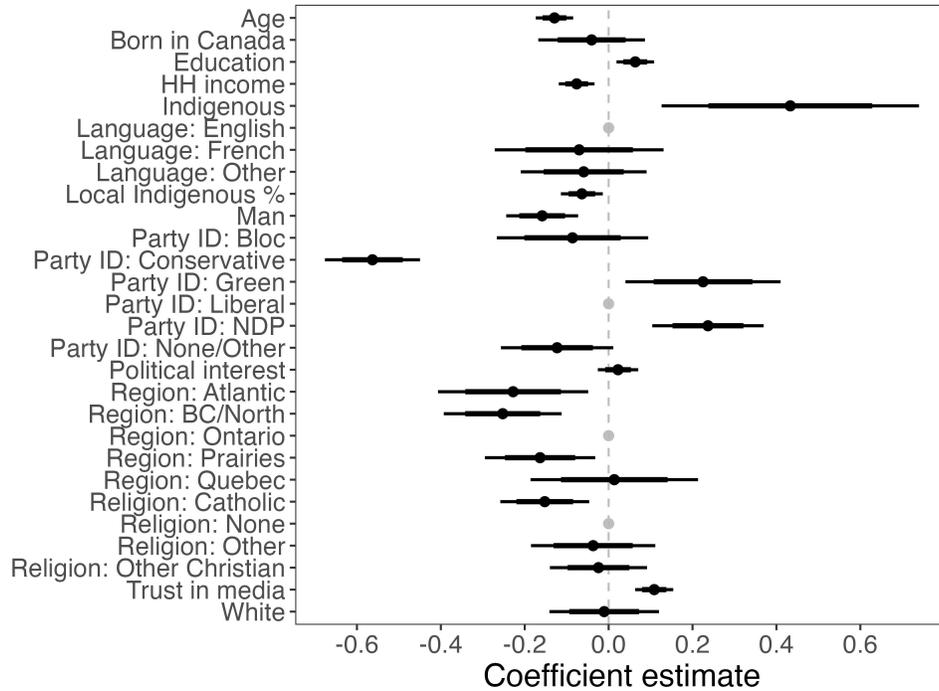


Figure A1: Correlates of belief in systemic racism

Using data from the 2019, 2020 and pre-discovery 2021 C-Dem Democracy Checkup surveys, this plot presents coefficient estimates from an OLS model regressing beliefs in systemic racism on the variables listed on the y -axis. Bars indicate standard HC2 95% and Bonferroni-adjusted 95% confidence intervals. The scale has been standardized so that estimates imply effects in terms of standard deviation changes, with higher values indicating stronger beliefs in the existence of systemic racism. All explanatory variables are binary except for Age, Education, Political Interest and HH income, which have been standardized such that the coefficient represents an implied effect of a one-standard deviation change. Reference categories for categorical variables are indicated by grey points.

A.3 Public mourning after the discoveries

After the first announcement of unmarked graves in Kamloops on May 27, 2021, Canadians across the country engaged in numerous public acts of mourning. Demonstrations, candlelight vigils, and remembrance walks saw hundreds of people attend in large cities, small towns and Indigenous communities. In many locales, children’s shoes were assembled as a memorial to the lives that were lost at residential schools.

To illustrate the magnitude of this public outpouring of grief, I recorded every mention of an event commemorating the deaths of Indigenous children in Canadian cities that occurred in the two weeks after the Kamloops announcement. Figure A2 summarizes the data. Events

were not specific to one region and took place throughout the week as the news gradually became more widely known. In total, I identified almost 90 events in this period, although the true number is likely higher because not all events were described in the media or online.

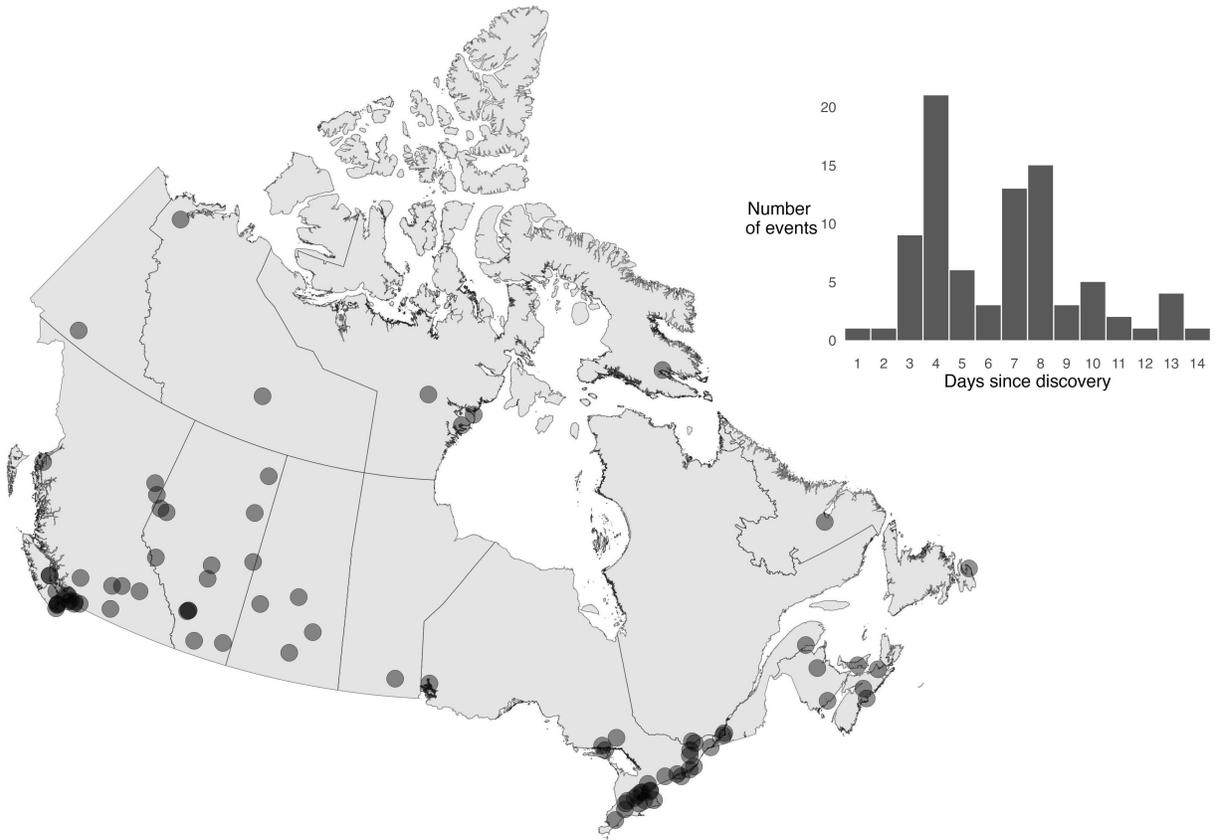


Figure A2: Location and frequency of vigils, May 28 to June 10

B. QUASI-EXPERIMENT

B.1 Balance checks

The fundamental assumption in my analysis of public opinion after the first discovery is that those who were surveyed just before the discovery was announced are similar on average to those surveyed immediately afterwards. To test this assumption, I first compare the means on a host of presumably stable covariates in Table A3. Overall, those in the pre- and post-discovery samples are remarkably similar on these observable dimensions, except those surveyed after the first discovery are about two years younger and have 0.4% fewer Indigenous people living in their electoral district.

I also conduct a separate check by regressing an indicator for being surveyed after the discovery on the same set of covariates to test whether the differences persist after conditioning on other possible sample differences. The standardized coefficient estimates predicting post-discovery status are summarized in Figure A3. The results in this test are qualitatively similar: few variables exhibit meaningful differences between the two samples. Only respondent age and the local Indigenous population percentage are statistically distinguishable from zero and even in those cases, the magnitudes are small: a one standard deviation increase in age and the local Indigenous percentage is associated with a 3 and 2 p.p. lower likelihood of being treated, respectively. The only other notably large coefficients are those associated with speaking French, living in Quebec and supporting the Bloc Québécois, but collinearity among these variables may be inflating the estimates given the balance reported in Table A3. Overall, these analyses suggest there are few discrepancies between the pre- and post-discovery samples and that they are small in size. Nonetheless, I control for all variables reported here in the ATE estimation.

Table A3: Sample characteristics by treatment status

	Average		
	Pre-discovery	Post-discovery	Difference
Man	0.50	0.47	0.03
Age	50.8	48.8	2.00*
White	0.80	0.79	0.01
Bachelor's degree	0.43	0.42	0.01
Household income	\$79,474	\$81,013	\$1,539
Catholic	0.28	0.29	0.01
Other Christian	0.23	0.23	0.00
Not religious	0.39	0.38	0.01
Born in Canada	0.80	0.81	0.01
Political interest (0 to 10)	6.46	6.34	0.12
Region: Ontario	0.39	0.41	0.02
Region: Quebec	0.26	0.25	0.01
Region: BC	0.10	0.11	0.01
Region: Atlantic	0.06	0.07	0.01
Local Indigenous %	0.04	0.04	0.00*
French-speaker	0.24	0.25	0.01
Party ID: Bloc	0.07	0.09	0.02
Party ID: Conservative	0.23	0.23	0.00
Party ID: Liberal	0.33	0.32	0.01
Party ID: NDP	0.14	0.14	0.00
Party ID: None/Other	0.22	0.23	0.01

*p<0.05 in *t*-test for difference-in-means.

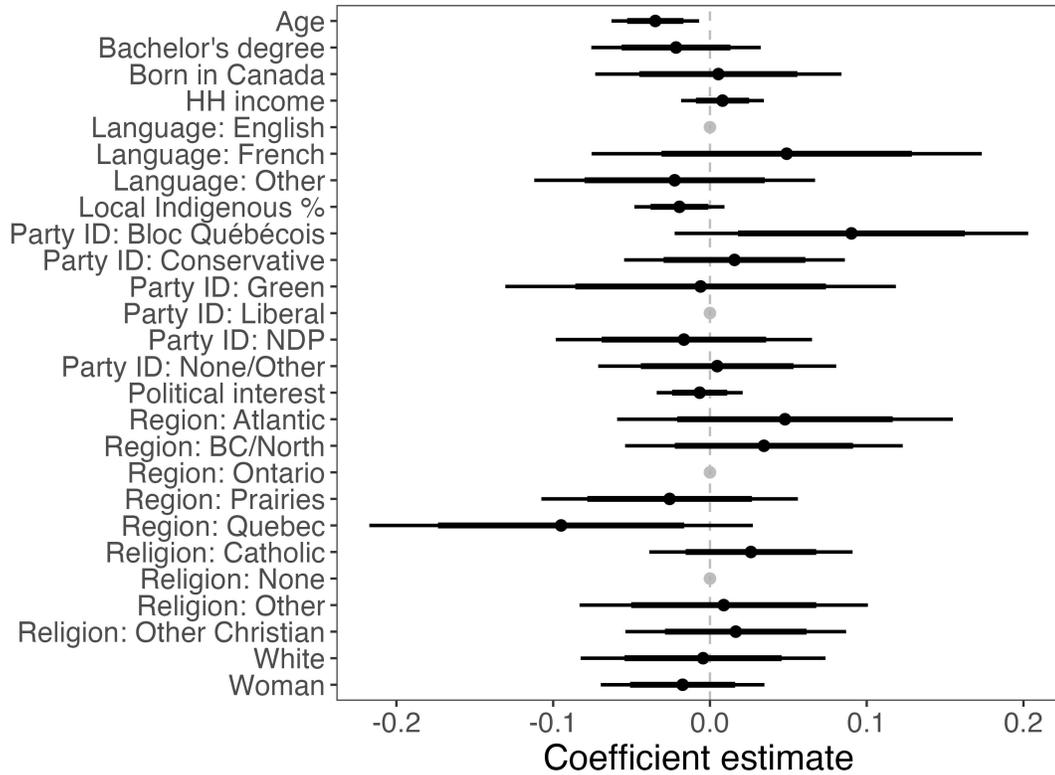


Figure A3: Quasi-experimental balance test

Plot presents coefficient estimates from an OLS model regressing treatment status on the variables listed on the y -axis. Bars indicate standard HC2 95% and Bonferroni-adjusted 95% confidence intervals. All variables are binary except for Age, Political Interest and HH income, which have been standardized such that the coefficient represents the implied effect of a one-standard deviation change. Reference categories for categorical variables are identified by grey points.

B.2 Complete covariate estimates

Table A4: Complete estimates for Table 1

		Belief in systemic racism
Intercept	0.000 (0.022)	-29.021* (10.440)
Surveyed after graves announcement	0.113* (0.032)	0.099* (0.030)
Party ID: Conservative		-0.534* (0.066)
Party ID: Green		0.308* (0.083)
Party ID: Liberal		0.122* (0.061)
Party ID: NDP		0.325* (0.068)
Party ID: None/Other		-0.113 (0.068)
Gender: Man		-0.262* (0.031)
Ethnicity: White		0.110* (0.047)
Household income		-0.000 (0.000)
Political interest		0.031* (0.007)
Born in Canada		0.007 (0.046)
religionNone		0.164* (0.038)
Religion: Other		0.004 (0.060)
Religion: Other Christian		0.066 (0.047)
Education		0.018* (0.009)
Language: French		-0.174* (0.078)
Language: Other/Multiple		-0.004 (0.052)
Local indigenous %		-1.738* (0.358)
Province: British Columbia		-0.046 (0.068)
Province: Manitoba		0.038 (0.095)
Province: New Brunswick		-0.197 (0.123)
Province: Newfoundland and Labrador		-0.099 (0.124)
Province: Nova Scotia		-0.081 (0.109)
Province: Ontario		0.119* (0.053)
Province: Prince Edward Island		0.024 (0.294)
Province: Quebec		0.214* (0.088)
Province: Saskatchewan		-0.150 (0.105)
Born 1930s		-0.161 (0.365)
Born 1950s		-0.358 (0.358)
Born 1950s		-0.552 (0.371)
Born 1960s		-0.692 (0.396)
Born 1970s		-0.810 (0.424)
Born 1980s		-0.862 (0.458)
Born 1990s		-0.821 (0.495)
Born 2000s		-0.711 (0.524)
Observations	3,849	3,752
Controls	No	Yes
R ²	0.003	0.197

Coefficients are expressed in terms of pre-announcement standard deviations. * p<0.05

B.3 Belief in systemic racism by survey date

Figure A4 charts the average agreement with the two systemic racism items by survey date. After the initial announcement on May 27, agreement trends steadily upward as the story became more widely known.

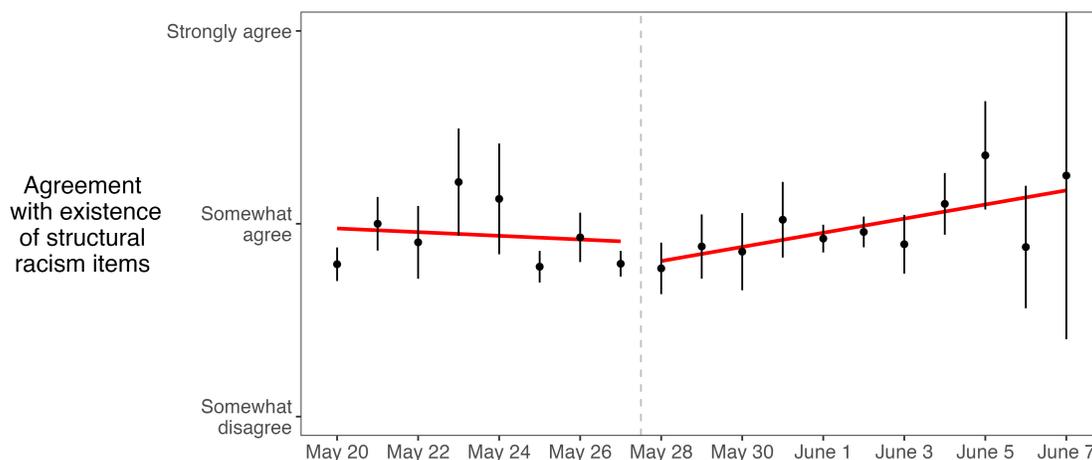


Figure A4: Agreement with systemic racism items by survey date

Plot presents mean and 95% confidence intervals for the average reported agreement with the existence of systemic racism among respondents each day the survey was in the field.

Partly this is because the news was first reported late on a Thursday on the west coast. Media did not begin covering the story intensely until after the weekend, mostly beginning on June 1 (see D.3). In Table A5, I re-estimate the models from the quasi-experiment using this date as the first true day of exposure to the story. Under this specification, the effects of being surveyed after the announcement became widely covered in the media are around 15 to 20% larger than the estimates reported in the main text.

Table A5: Effects of unmarked graves news using alternative exposure date

	Belief in systemic racism	
Surveyed after May 31	0.136* (0.034)	0.116* (0.031)
Observations	3,849	3,752
Controls	No	Yes
R ²	0.004	0.196

Coefficients are expressed in terms of pre-announcement standard deviations. In model 2, the following covariates are included but not reported: gender, born in Canada, education, household income, party ID, political interest, religion, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. *p<0.05

B.4 Effects on individual survey items

Table A6: Unmarked graves discovery and individual systemic racism items

	Belief in systemic racism			
	Colonialism item		Deservingness item	
Surveyed after graves discovery	0.098* (0.030)	0.089* (0.028)	0.103* (0.032)	0.088* (0.030)
Observations	3,852	3,755	3,850	3,753
Controls	No	Yes	No	Yes
R ²	0.003	0.175	0.003	0.166

Coefficients are expressed in terms of pre-discovery standard deviations. In models 2 and 4, the following covariates are included but not reported: gender, born in Canada, education, household income, party ID, political interest, religion, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. *p<0.05

B.5 Effects on racial minorities feeling thermometer

The main analysis focuses on respondents’ beliefs about the existence of anti-Indigenous systemic racism, rather than, for example, affective feelings toward the outgroup. To get at this latter outcome, I re-analyze the quasi-experiment with a 0 to 100 feeling thermometer score toward “racial minorities.” Table A7 displays the coefficient estimates. Whether adjusting for covariates or not, neither of the estimates are statistically significant and both are only a fraction of the size of the estimate for systemic racism attitudes (~ 0.1 standard deviations).

That being said, this is not a definitive test for whether the grave discovery moved affective attitudes toward the victim group. Respondents were asked about racial minorities in general, rather than Indigenous Peoples in particular. While some may have conflated the two groups in their responses, it is impossible to know how prevalent this pattern was in retrospect. For this reason, we cannot rule out the possibility that outgroup attitudes changed beyond the observed shifts in systemic racism beliefs reported in the main analysis.

Table A7: Unmarked graves discovery and feeling thermometer toward racial minorities

	Racial minorities feeling thermometer	
Surveyed after graves discovery	0.045 (0.033)	0.029 (0.031)
Observations	3,772	3,681
Controls	No	Yes
R ²	0.001	0.143

Coefficients are expressed in terms of pre-discovery standard deviations. In model 2, the following covariates are included but not reported: gender, born in Canada, education, household income, party ID, political interest, religion, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. *p<0.05

B.6 Placebo test: Attitudes toward immigrants

Table A8: Placebo test: Unmarked graves discovery and attitudes towards immigrants

	Anti-immigrant attitudes	
Surveyed after graves discovery	0.010 (0.033)	0.003 (0.031)
Observations	3,851	3,755
Controls	No	Yes
R ²	0.000	0.186

Outcome is average agreement with the following two items: “Too many recent immigrants just don’t want to fit in to Canadian society” and “Immigrants take jobs away from other Canadians.” Coefficients are expressed in terms of pre-discovery standard deviations. In model 2, the following covariates are included but not reported: gender, born in Canada, education, household income, party ID, political interest, religion, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. *p<0.05

C. OVER-TIME ANALYSES

C.1 Sample details

In investigating attitudinal persistence, I rely on responses to the Indigenous resentment items in the 2019, 2020, 2021 and 2022 C-Dem Democracy Checkup surveys and the 2021 Canadian Election Study. The sampling strategy is nearly identical in all four surveys, relying on online recruitment and quotas for age, gender, province and language (in Quebec), based on the 2016 Canadian Census. In the 2021 CES, the data includes “oversampled” responses that collected overflow responses from full quotas and in the 2022 Democracy Checkup, an oversample of Quebec respondents was conducted; I exclude both sets of oversample responses in my analyses. Respondents needed to be 18 years of age or older, and Canadian citizens or permanent residents to participate. In all analyses, responses are unweighted.

One noteworthy change in sampling is that for the 2019 and 2020 surveys, C-Dem contracted with Dynata to recruit respondents, while Leger is used from 2021 onwards. The overall level of belief in systemic racism is lower in the Dynata samples, although this is confounded by time. Importantly, the main pattern of interest – a short-term decrease in resentment in the 2021 Democracy Checkup followed by an increase in the resentment months later in the 2021 CES – is observable strictly within the surveys fielded by Leger. These facts help ameliorate any concerns that the over time trends I observe are due to changes in sample characteristics, although I investigate this possibility in more detail below.

C.2 Changes in sample composition over time

One explanation for the changes we observe in beliefs about systemic racism over the last few years – and in particular the reversion in attitudes after the unmarked graves discoveries in 2021 – is that the samples of respondents are changing rather than just their reported attitudes. However, the survey is a general omnibus political attitudes poll, not one focused exclusively on intergroup attitudes, so it is unlikely that there would be response bias specifically on this issue after the unmarked graves announcements. Nonetheless, to evaluate the extent of changing sample characteristics, I estimate a multinomial logistic regression where the outcome is a categorical variable indicating the survey wave a respondent belongs to and the predictor variables are demographic characteristics. If there are no differences in respondent characteristics across survey waves, the variables should not predict the survey that each respondent comes from. Note that in the Summer 2019 survey, there is no question that identified respondents' race, so this survey is not used in the analyses here. As the plot in the main text showed, however, there was little change in respondents' attitudes between 2019 and 2020.

Table A9 summarizes the model; note that the reference category for the outcome variable is the May 2021 (the field dates of which included the first unmarked graves discovery). The coefficient estimates represent the change in log odds of appearing in the survey in the

columns relative to the May 2021 survey given a unit change in the predictor variables while all other predictors are held constant. The model reveals some notable changes between survey waves: for example, respondents generally became younger over each successive wave and, relative to those in the May 2021 survey, those in later surveys were more likely to be born in Canada, have higher incomes, and less likely to be white.

Given these relevant differences in sample characteristics, I attempt to control for all changes over time in respondents' observable characteristics. Specifically, I estimate an OLS model predicting systemic racism beliefs across all survey waves (except 2019; see above) based on the variables in Table A9 and then summarize the model residuals across each survey. Figure A5 presents the results. This plot captures the average belief in systemic racism in each survey after partialing out observable variables. Note that larger residuals indicate a greater belief in systemic racism than would have been predicted by all of the time-invariant pre-treatment covariates. Encouragingly, the plot shows the same pattern as the raw scores in the main analysis: belief in systemic racism becomes stronger after the initial unmarked grave discovery (marked by a dashed vertical line), but return to baseline levels in the later surveys. These results suggest that the reversion in attitudes is not driven simply by a change in sample characteristics over time.

Table A9: Predicting sample membership
from respondent characteristics

	Survey wave (Reference category is May 2021)		
	May 2020	September 2021	May 2022
Age	-0.086*	0.065	0.222*
	(0.023)	(0.022)	(0.020)
Man	0.087	-0.068	-0.089
	(0.044)	(0.043)	(0.039)
Born in Canada	0.071	0.199*	0.749*
	(0.062)	(0.061)	(0.057)
Region: BC/North	0.271	0.290	-0.012
	(0.103)	(0.101)	(0.093)
Region: Ontario	0.002	0.037	-0.206
	(0.089)	(0.088)	(0.080)
Region: Prairies	-0.010	0.067	-0.012
	(0.098)	(0.097)	(0.088)
Region: Quebec	-0.248	-0.177	0.084
	(0.120)	(0.117)	(0.105)
Religion: None	-0.102	0.175*	0.126
	(0.054)	(0.053)	(0.048)
Religion: Other	-0.027	0.097	0.056
	(0.081)	(0.079)	(0.073)
Religion: Other Christian	-0.069	-0.063	-0.010
	(0.063)	(0.063)	(0.057)
Party ID: Conservative	0.100	0.190	0.043
	(0.099)	(0.098)	(0.086)
Party ID: Green	0.063	-0.188	-0.457*
	(0.128)	(0.131)	(0.117)
Party ID: Liberal	0.031	0.008	-0.074
	(0.094)	(0.093)	(0.081)
Party ID: NDP	-0.169	0.310	0.162
	(0.106)	(0.102)	(0.091)
Party ID: None/Other	-0.214	0.262	-0.372*
	(0.103)	(0.099)	(0.089)
Household income	0.006	0.083*	0.066*
	(0.021)	(0.020)	(0.019)
Local Indigenous %	0.015	0.020	0.035
	(0.023)	(0.022)	(0.020)
Education	-0.047	-0.025	0.048
	(0.022)	(0.022)	(0.020)
Political interest	0.070*	-0.023	0.005
	(0.023)	(0.022)	(0.020)
White	0.076	-0.199*	-0.148
	(0.061)	(0.059)	(0.055)
Language: French	0.109	0.206	0.119
	(0.097)	(0.095)	(0.085)
Language: Other/Multiple	-0.084	-0.056	1.217*
	(0.074)	(0.071)	(0.063)
Akaike Inf. Crit.	80,011.560		

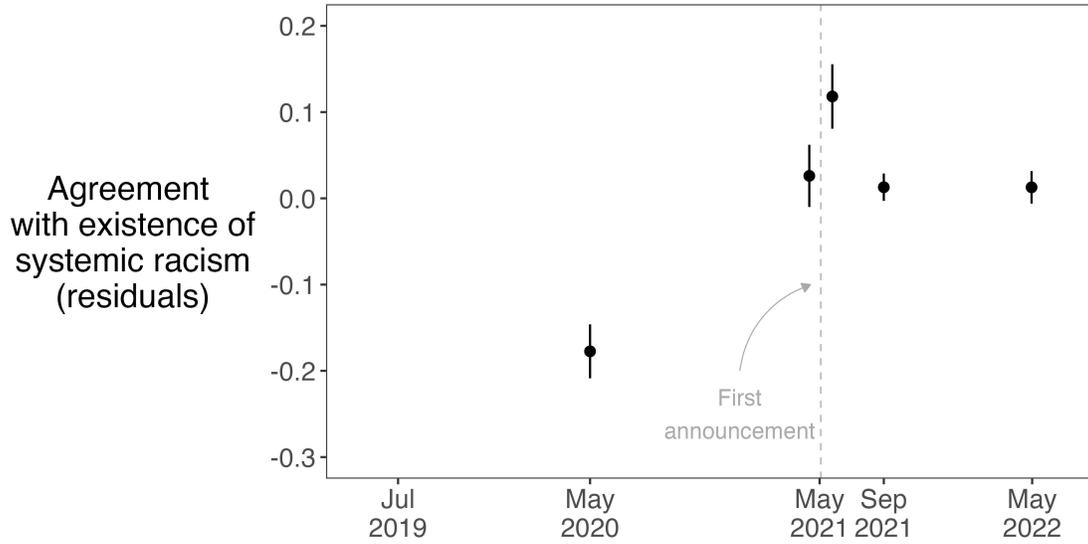


Figure A5: Residualized beliefs in systemic racism, 2019 to 2022

Plot presents average and 95% confidence intervals for respondents' residualized systemic racism scores in each survey wave. Residuals calculated from an OLS model including the following predictors: age, gender, white, born in Canada, region, religion, party ID, household income, language, education and political interest. The vertical dashed line indicates the initial discovery of unmarked graves in 2021.

C.3 Systemic racism items separately over time

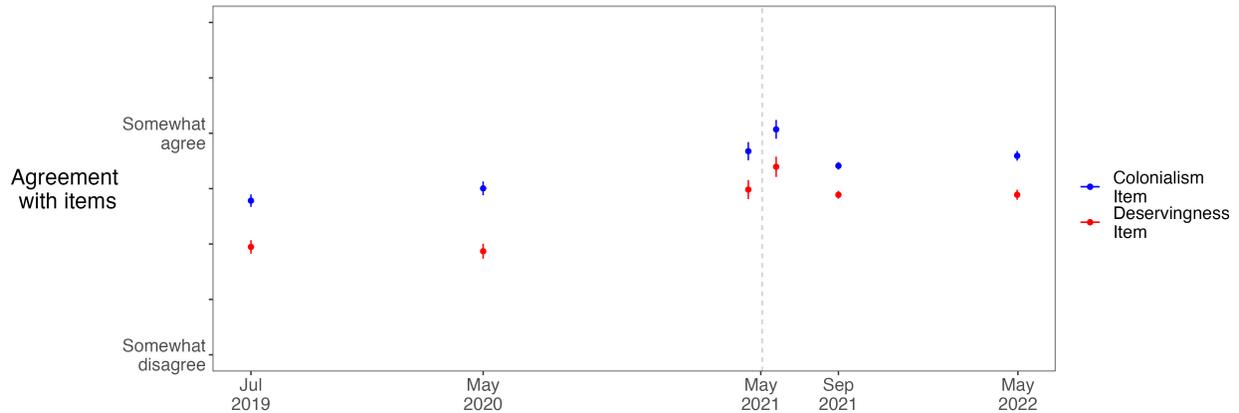


Figure A6: Beliefs in systemic racism, 2019 to 2022

Plot presents average and 95% confidence intervals for the mean for each of the two items measuring beliefs in systemic racism. In all but the September 2021 survey, this question was asked on a four-point Likert scale. For that specific survey, a five-point scale was used, but responses have been rescaled to match the four-point scale (see Appendix Figure A7 for individual response level prevalence over time).

C.4 Response-level frequency over time

The over time comparisons in the main text rely on five different surveys. In all but one of those surveys, the items asking about agreement with the existence of systemic racism are measured on a four-point scale. In the main text, responses from the only survey using a five-point scale (the September 2021 survey) are rescaled to match the other years’ data.

In Figure A7, I show that the same substantive conclusions can still be drawn without rescaling the data. The two plots present that proportion of respondents answering with each response level to each systemic racism item in each survey wave. When the grave announcement was first made (indicated by the gray dashed line), all response categories became less prevalent except for the one that indicated strong agreement with the existence of systemic racism. In the September 2021 survey, we again see nearly all response categories becoming less prevalent, but this time because a “neutral” option has been introduced. That being said, more respondents sorted out of the “strongly agree” option than the “strongly disagree” option between the two 2021 waves, indicating a reversion in attitudes. Finally, when the “neutral” option is again removed in 2022, we do not see many respondents returning to a “strongly agree” position, suggesting possibly that attitudes stabilized after 2021.

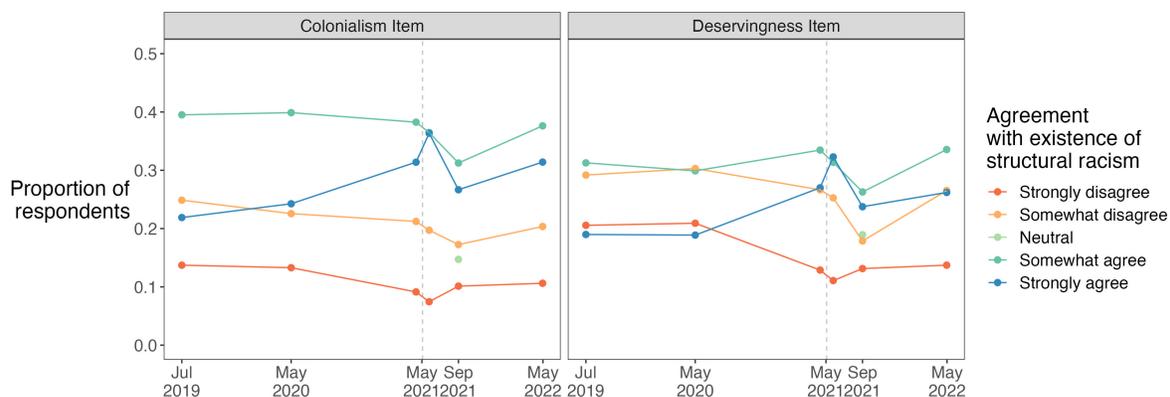


Figure A7: Beliefs in systemic racism by response level, 2019 to 2022

Plot presents the proportion of respondents offering each response level to the two systemic racism items in each survey wave. The colour scale has been defined so that higher values indicate greater agreement with the existence of systemic racism. Except for the September 2021 survey, these items was asked on a four-point Likert scale.

C.5 Indigenous feeling thermometer over time

The main analyses in this study focus on beliefs in systemic racism. In Figure A8 I instead look at a measure of affective attitudes toward Indigenous people over time: non-Indigenous people's average responses when asked to rate how they feel about Indigenous Peoples on a scale from 0 to 100, with larger values indicating more positive views. While this plot only allows for a descriptive look at attitudes over time, one pattern worth noting is that the average score in 2021 after the grave discoveries does not appear to be much different than responses among online respondents since 2015. There is no increase in favourable attitudes as a result of the revelation of historical injustices, further corroborating the main analyses' finding of little long-term attitudinal change.

This data also reveals an important mode effect. In 2015, the Canadian Election Study was fielded both online and over the phone. Those completing the survey online rated Indigenous Peoples nearly 11 points lower on the feeling thermometer than those who spoke to an enumerator over the phone ($p < 0.001$). These results suggest that online surveying may be significantly reducing respondent incentives to provide socially desirable responses.

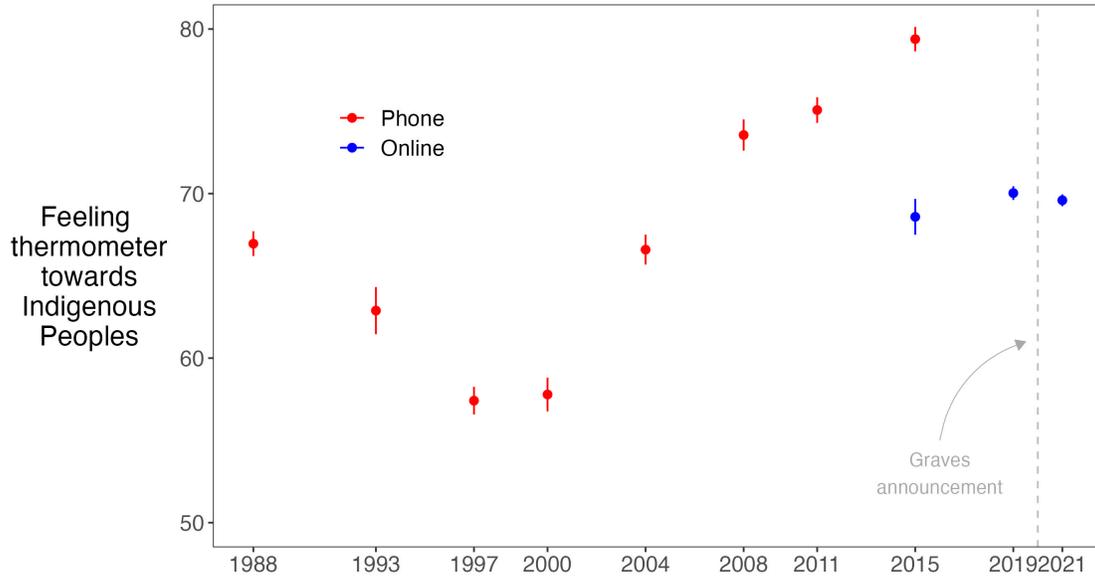


Figure A8: Feeling thermometer toward Indigenous Peoples, 1988 to 2021

Plot presents the average and 95% confidence intervals for the feeling thermometer score in each survey year. The specific language of the item changed over time: respondents gave their feelings towards “Native peoples” in 1988, “Aboriginal peoples” from 1993 to 2015, and “Indigenous Peoples” from 2019 onwards. In 2019, point represents the pooled average within two surveys: the Canadian Election Study and the Democracy Checkup.

C.6 Racial minorities feeling thermometer over time

In the main text, I discuss the murder of George Floyd in the United States and the subsequent protests that occurred both in that country and in Canada as one possible explanation for the increase in respondents’ willingness to acknowledge anti-Indigenous systemic racism in the period *before* the graves were discovered. To shed light on this possibility, I report on changes in non-Indigenous Canadians’ attitudes toward “racial minorities” using a feeling thermometer score in Figure A9. (This figure complements the analyses in B.5, which investigates the grave discovery’s effects on this same variable).

Paralleling the main text findings with regard to beliefs in systemic racism (see Figure 3), there is a simultaneous improvement in attitudes toward racial minorities between May 2020 (just before George Floyd’s murder) and May 2021, before the first discovery was announced. Feeling thermometer scores improve by 5.1% between these dates, while acknowledgement

of anti-Indigenous systemic racism increases by 8.6%. Of course, we should be cautious in drawing conclusions from this analysis because the feeling thermometer is measuring a different underlying concept, and not a different racial group. For this reason, I cannot rule out alternative explanations for the change in pre-treatment attitudes, but given the lack of a trend in systemic racism beliefs between the summer 2019 and May 2020 surveys, George Floyd’s murder seems like a plausible explanation.

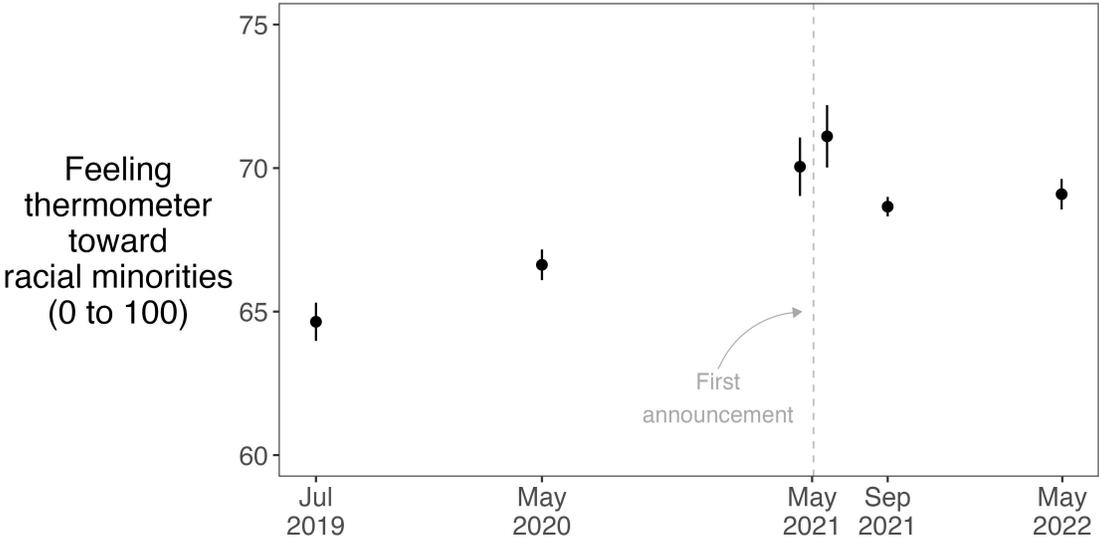


Figure A9: Feeling thermometer toward racial minorities, 2019 to 2022

Plot presents average and 95% confidence intervals for the feeling thermometer score reported by respondents towards racial minorities. The vertical dashed line indicates the initial discovery of unmarked graves in 2021.

C.7 Public prioritization of Indigenous issues over time

The main analysis in this paper looks at changes in beliefs about systemic racism. Yet the public’s interest in Indigenous policy issues exhibited a similar pattern over time. I collated data from the Angus Reid Institute, a non-profit, non-partisan research foundation, on the percentage of respondents listing “Indigenous issues” or “reconciliation” among the top three policy issues they care about most when asked in repeated cross-sectional surveys from the past three years.

As Figure A10 shows, around 10% of Canadians provided this response in the months

just before the first discovery in May 2021. Immediately after the unmarked grave announcements, this proportion more than doubled. Attention to Indigenous issues had not been as high since early 2020, when Indigenous communities protested against the construction of a natural gas pipeline in British Columbia. In that case, however, interest in Indigenous issues was less related to historical injustices than to conflicts over land rights and resource development.

Since the policy interest variable can reflect a diversity of viewpoints toward Indigenous Peoples, it cannot be used to infer opinion change. However, Angus Reid conducts polls at a higher frequency than the surveys used in the main analysis, so it is useful to examine how attention changed over the study period. After the initial increase in prioritization of Indigenous issues, interest returned to baseline levels within a year and, over the following months, trended even further downward. These findings are generally similar to the results for beliefs in systemic racism. And as with that outcome, the changes are similar across partisan groups: despite significant pre-discovery differences in the importance that supporters of different parties attached to Indigenous issues, the interest levels of each group increased only temporarily before quickly decaying in the months after the discoveries.

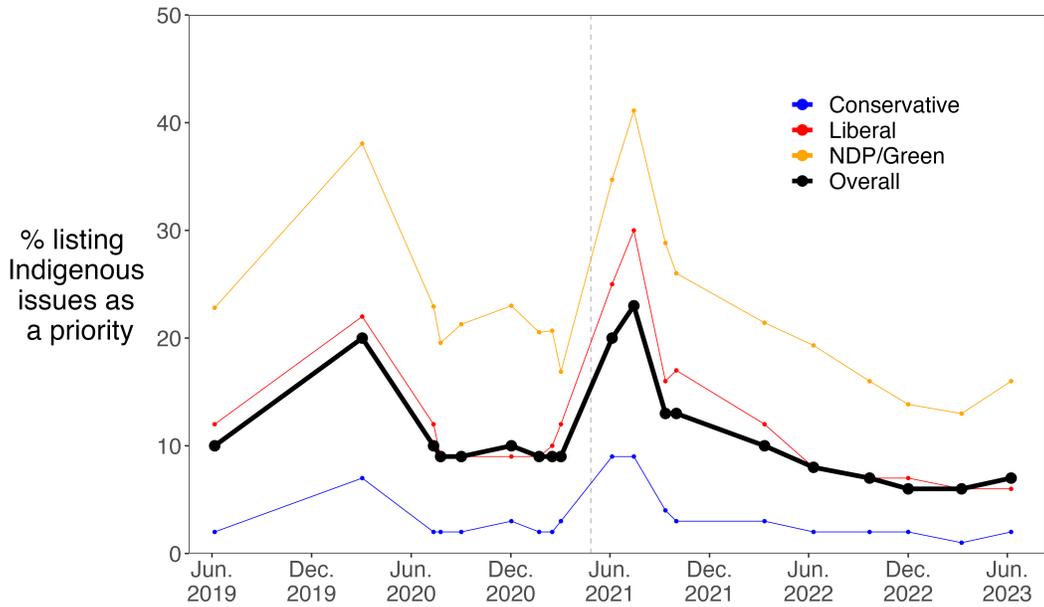


Figure A10: Public attention to Indigenous issues, July 2020 to September 2022

Data are from Angus Reid Institute public opinion polls; the *y*-axis indicates the percentage of respondents, by party and overall, that selected “Indigenous issues” as one of their top-three choices when responding to the question “Thinking of the various issues facing Canada today, which ones do you personally care about the most?”

D. MEDIA ANALYSIS

D.1 Data sources and text pre-processing

The media content analysis in this study is based on all articles published in Canada’s six largest English-language newspapers and three major regional newspapers between January 1 and December 31, 2021. The estimated political slant of each outlet is presented in Figure A11, according to data from Media Bias/Fact Check.

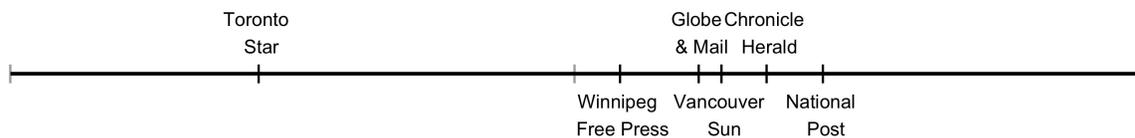


Figure A11: Media source biases

The article data were downloaded as a document-frequency matrix from ProQuest’s database, which does not carry French-language Canadian newspapers.

The following pre-processing steps were carried out before estimating topic models:

1. Terms were tokenized into unigrams and converted to lowercase.
2. Stop-words, punctuation, numbers and terms appearing in less than 1% of articles were removed.
3. Articles shorter than 100 words were removed.
4. Duplicate articles were removed in two ways:
 - (a) Duplicate stories appearing in the same outlet on the same date with the same title were removed.
 - (b) A cosine distance matrix was calculated measuring the similarity between each article in the dataset with all other articles. If articles had a cosine similarity score greater than 0.9, a random article among the similar articles was chosen to remain in the dataset and others were removed. This step is necessary because several of the outlets share the same parent company and publish syndicated articles.

After these steps, the corpus comprised 81,544 articles. Note that terms were not stemmed or lemmatized, following topic modelling advice in Schofield and Mimno (2016).

D.2 Topic model estimation and validation

I begin my analysis of media coverage by training a Latent Dirichlet Allocation (LDA) topic model. While the estimation of an LDA topic model is straightforward, one challenge for the analyst is selecting the parameter K , or the number of topics to be identified by the model. This task requires validation based on both statistical and interpretability considerations (Grimmer and Stewart 2013). In testing various values of K , a common problem was that the “residential schools” topic would become amalgamated with a broader “Indigenous” topic for all $K < 150$. Given my interest in the residential schools topic specifically, I therefore considered only values above this threshold. To select the optimal K beyond this criterion, I estimated LDA models for $K = \{150, 200, 250, 300\}$.

In general, there is little disagreement among these models about the prevalence of the residential school topic in each article. The table below presents the correlation between the estimated residential school topic prevalence across different values of K :

	LDA K			
	150	200	250	300
150	1.00			
200	0.93	1.00		
250	0.93	0.97	1.00	
300	0.92	0.97	0.97	1.00

The correlations are extremely high, scoring between 0.92 and 0.97. The top terms for each residential school topic are also very similar. In this context, the choice of K will likely make little substantive difference. Similarly, the log-likelihood is not significantly different across specifications. Given these considerations, I set $K = 250$. I cannot claim that this is the definitive number of topics in the corpus, but the results from this specification appear substantively meaningful based on the output.

To illustrate this, Figure A12 plots the terms most associated with the residential school topic, alongside the general Indigenous topic for comparison. The residential school topic

picks up on coverage related to the schools themselves, but also the graves, the Truth and Reconciliation Commission, the debates over history, and also discussions of commemorating Canada Day and the National Day for Truth and Reconciliation. The general Indigenous topic is different; here the most important terms deal with Indigenous communities and their land and governance systems.

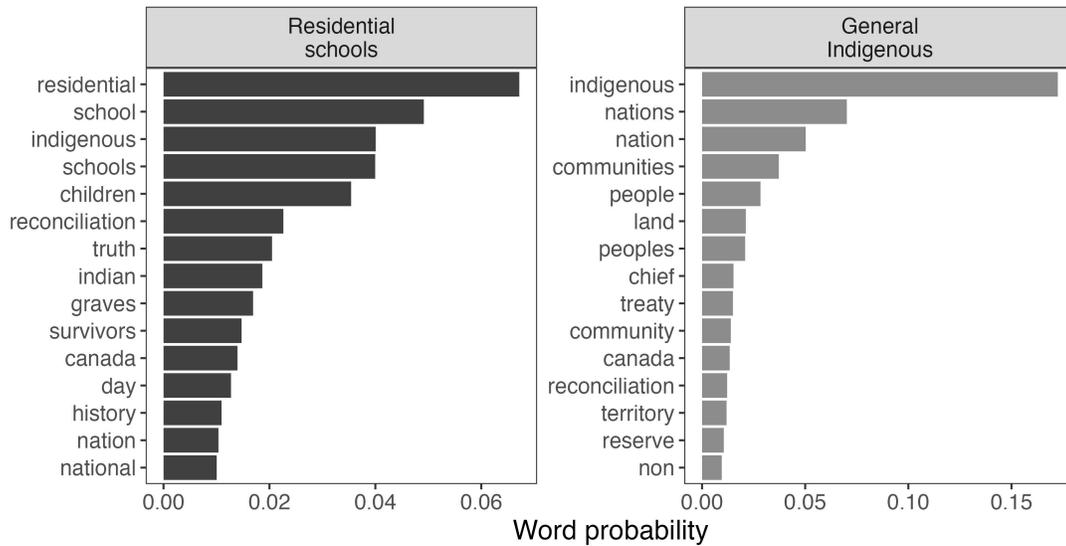


Figure A12: Terms most associated with Indigenous topics in LDA model estimates

Plot presents the fifteen terms that are most associated with the residential schools and general Indigenous topics in the LDA model along with their word probabilities for those topics.

Table A10 presents the documents in the corpus most associated with the residential school and general Indigenous topics. The articles related to residential schools have strong face validity: all of these stories deal with the residential schools or the announcements of unmarked graves. By contrast, the articles related to the general Indigenous topic are of a different nature, dealing with issues of land rights and Indigenous governance, with no mention of the residential schools topic. In each set of articles, the topic prevalence estimates for the opposite topic are very small. That is, the model predicts that the articles most related to residential schools are not strongly related to the general Indigenous topic. These results provide additional confidence that the residential schools topic is coherent and capturing my specific focus on this historical injustice.

Table A10: Top documents per topic

<i>Panel A: Top documents related to residential schools topic</i>				
Source	Date	Title	RS topic	Ind. topic
Vancouver Sun	Jun. 25	'We are treating this like a crime scene'; Saskatchewan first nation says it found 751 unmarked graves at residential school site	0.50	0.00
Chronicle - Herald	Jun. 25	Saskatchewan site likely contains 751 unmarked graves	0.50	0.01
Vancouver Sun	May 28	Remains of 215 children found at residential school; Kamloops First Nation says discovery confirms what community had known	0.47	0.01
Globe and Mail	May 29	Remains found at B.C. residential school site	0.47	0.02
Toronto Star	Jul. 16	First Nation vows it will ID Kamloops remains	0.46	0.00
Globe and Mail	Oct. 1	Day to reflect and remember	0.46	0.04
Winnipeg Free Press	Jun. 7	The children can no longer be ignored	0.46	0.00
Globe and Mail	Jun. 17	Well-wishers overwhelm First Nation	0.46	0.02
Globe and Mail	Aug. 11	First Nations begin process to find children who didn't return from Vancouver residential school	0.46	0.00
Toronto Star	Jun. 1	More Canadian sites believed to hold unmarked graves	0.45	0.00
<i>Panel B: Top documents related to general Indigenous topic</i>				
Source	Date	Title	RS topic	Ind. topic
National Post	Nov. 11	Learn before you blockade	0.00	0.35
Winnipeg Free Press	Dec. 1	At last, a legislative land acknowledgment	0.01	0.33
Toronto Star	Dec. 5	What would giving land back look like?	0.00	0.32
Winnipeg Free Press	Jun. 28	Treaty No. 1 document can't convey full story	0.00	0.30
Globe and Mail	Feb. 23	A year after Wet'suwet'en crisis, questions over self-governance have only grown louder	0.00	0.30
Winnipeg Free Press	Jul. 5	Visionary Treaty 1 chiefs experienced strategists	0.00	0.30
National Post	Dec. 7	We're Wet' suwet' en, but Coastal GasLink protesters don't represent us	0.00	0.29
Winnipeg Free Press	Mar. 30	Circular fight for Indigenous rights	0.00	0.28
Vancouver Sun	Apr. 22	Paying dearly for 'unity'	0.00	0.28
Winnipeg Free Press	Aug. 1	Sesquicentennial of Treaty No. 1 signing honoured	0.10	0.27

D.3 Residential school topic prevalence during quasi-experiment

Using the estimates from the LDA model described in Appendix D.2, Figure A13 presents the proportion of newspaper coverage related to the residential schools topic by survey date during the quasi-experiment. After the initial announcement of suspected unmarked graves was made on Thursday May 27, coverage steadily increased over the weekend before jumping on June 1. In Appendix B.3 I re-estimate my main models using this alternative date as the onset of exposure to the residential schools information.

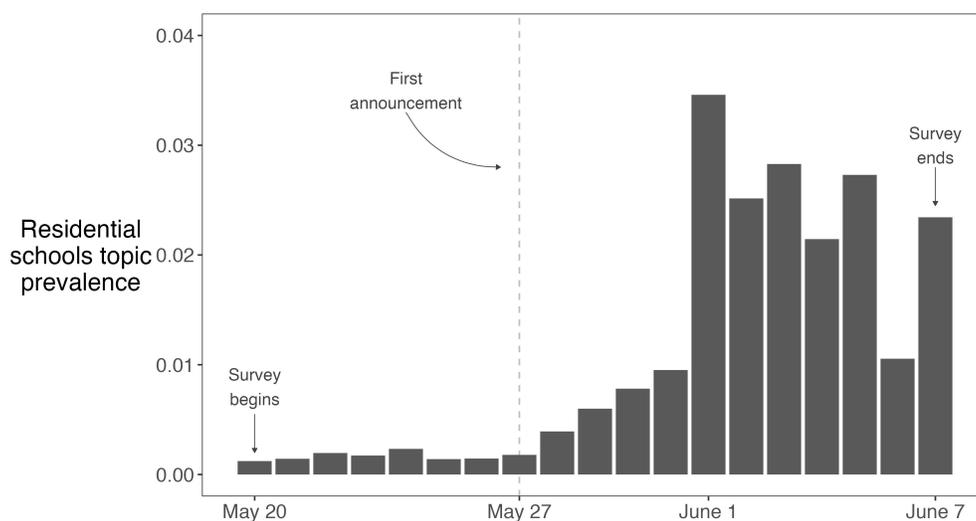


Figure A13: Prevalence of residential school topic during quasi-experiment survey dates. Plot presents the daily average of the residential schools topic prevalence estimated from an LDA topic model on the full-text of every news story in six of Canada’s largest English-language newspapers.

D.4 Residential schools and unmarked graves media coverage, 2000 to 2022

In the main text, I present the results of a topic model focusing on the prevalence of coverage related to residential schools during 2021. To look at a longer time horizon, I assembled a separate corpus of every article published in Canada’s three largest English-language newspapers between 2000 and 2022.

Figure A14 charts the percentage of news stories each month that mention the phrase “residential school(s).” The plot reveals that Canadians were exposed to the residential

school history more intensely after the unmarked grave announcements than at any point in the previous two decades. While coverage increased after the TRC issued its final report in 2015, no other period comes close to the volume of articles referencing this history in the summer of 2021.

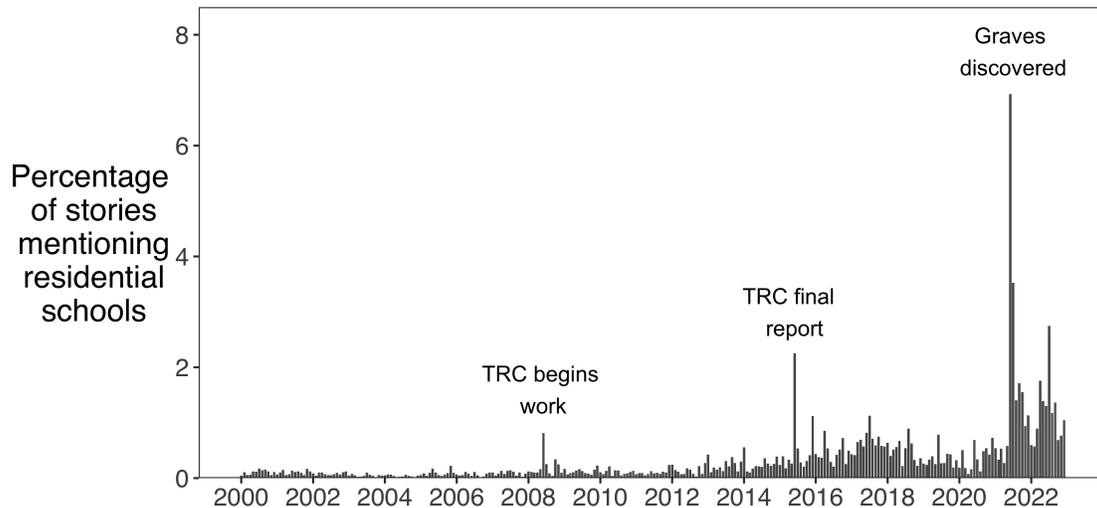


Figure A14: Percentage of newspaper stories mentioning residential schools, 2000 to 2022
Plot presents the percentage of all new stories that mention “residential school(s)” in the *Globe and Mail*, *National Post* and *Toronto Star* by month.

That being said, Figure A14 does show that the residential school history was covered to some degree before the events of 2021. However, the news about unmarked graves represented an entirely new dimension to the story, of which most Canadians were unaware. While the Truth and Reconciliation Commission signalled that graves likely existed at former schools and recommended further action to identify and protect these sites (i.e. Calls to Action 73 to 76), this information did not appear much in the media before 2021. In Figure A15, I count the number of articles each year that include the words “graves” and “residential school(s)” together in Canada’s largest English-language newspapers. Before 2021, there were essentially no articles discussing this topic, while that number increases dramatically after the Kamloops announcement. Moreover, many of the grave-related residential school articles before 2021 are likely false positives. I hand-coded 20 of the pre-2021

articles mentioning these two terms together and only three of them actually had to do with burials at the schools, while many were accidentally flagged because of phrases like “grave doubts.”

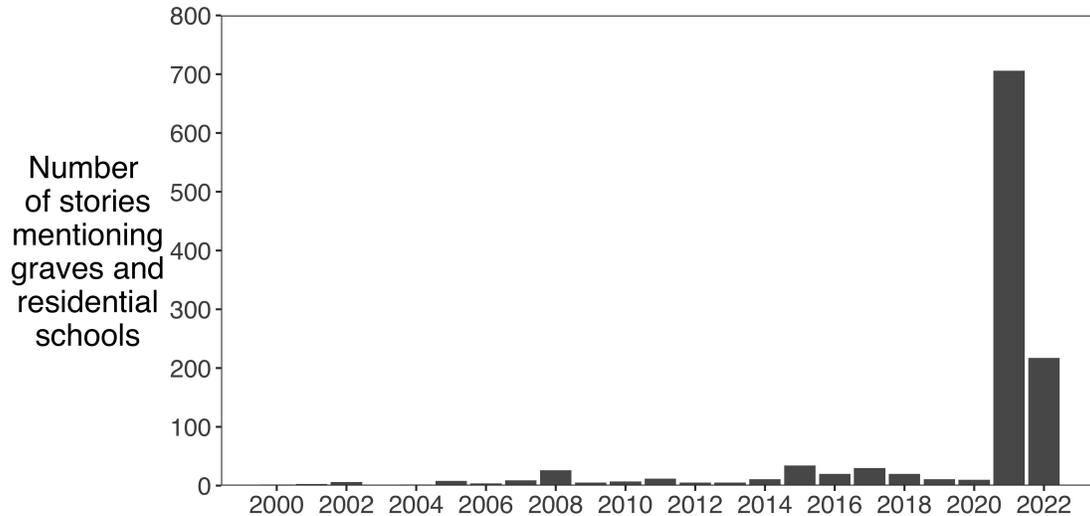


Figure A15: Number of articles mentioning “graves” and “residential schools,” 2000 to 2022
 Plot presents the number of all new stories that mention “residential school(s)” and “grave(s)” in the *Globe and Mail*, *National Post* and *Toronto Star* by year.

D.5 Framing analysis

To conduct my analysis of how the media framed the residential schools topic, I need to first identify the articles that are related to this topic. I relied on two-step classification process. First, I selected all articles containing the word “Indigenous” and appearing in the six largest English-language newspapers between the first announcement of unmarked graves in May 2021 and the end of the second survey in October. Within this set, I then read each article title and text to determine whether it was related to residential schools based on whether it was primarily concerned with:

- The history or legacies of residential schools
- The Truth and Reconciliation Commission
- Reconciliation for the residential school system
- Personal experiences in the schools

- Residential school denialism
- Policy and symbolic responses to the residential schools issue
- Church responses to the residential schools issue
- Non-indigenous reactions to the residential school history

If an article covered one of the above topics, I used it in the framing analysis. I did not flag articles as related to residential schools if they were about (i) Indigenous topics not directly related to residential schools, (ii) discrimination and racism against Indigenous people in general, or (iii) Indigenous leaders, community members or community initiatives unrelated to residential schools.

Initially filtering articles based on the appearance of the “Indigenous” term resulted in fewer false negatives than more restrictive decision rules (e.g. only selecting articles containing *both* “Indigenous” and “residential”) and machine learning algorithms. To illustrate the performance of this decision rule, I hand-coded 850 articles that both included and excluded the term “Indigenous” based on whether they related to the residential schools topic. While only 29% of the articles containing the word “Indigenous” were actually related to residential schools, none of the articles that did not mention this word were related to residential schools. In short, this decision rule casts a very wide net, identifying a number of articles that will not ultimately be used in the framing analysis, but it ensures that I have truly located all articles relevant to the topic of interest.

In total, 628 articles were related to residential schools and used in the framing analysis. To classify the framing of each article, I focus on the main way in which the residential schools issue is discussed in the article. Building on a classification developed by Benford and Snow, I consider whether the following frames are present:

- Descriptive: identifies a problem or injustice
- Solutions-oriented:
 - (a) identifies a solution to address an injustice
 - (b) attributes responsibility to actor for an injustice or its remedy
 - (c) makes a call to action or seeks to mobilize support for a solution

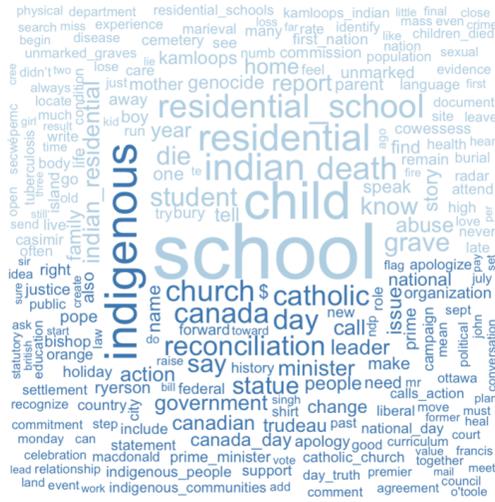
While an article may contain multiple frames (i.e. a “mixed” framing), for a frame to be coded as present it must be a central focus of the article and not appear tangentially. I exclude the $\approx 2.5\%$ of articles that do not contain either frame. To facilitate replicability, I select a passage that is representative of each frame identified in an article. Those quotations are available in the replication data along with article IDs that can be linked to the full text of each article in ProQuest’s database.

To validate the coding, I look at word usage in the articles containing each frame. Figure A16 is a wordcloud showing the most common terms in each set of articles. Figure A17 presents keyness statistics summarizing the terms that are most distinctive to articles with each frame. Together these plots show that, first, descriptive articles focus much more on the tragedy of the schools. They discuss the deaths, abuse and poor living conditions in these institutions as well as the individual victims. Importantly, they also talk more about the injustice itself: “residential school” and its related terms are much more common in descriptive coverage. By contrast, the solutions-oriented frames focus on responsibility – mentioning actors like the Catholic Church, the Pope, and Prime Minister Trudeau – and steps that can be taken to address the injustice. These articles are much more likely to discuss statue removals, cancelling Canada Day, spending money (“\$”) and taking action toward reconciliation.

In addition to the overall classification of articles as descriptive versus solutions-oriented, I also identify the topic of interest within each of these broader frames:

1. Descriptive
 - (a) Unmarked graves and deaths at the schools
 - (b) Assimilation and experiences in the schools
 - (c) Legacies of the school system
 - (d) Other
2. Solutions-oriented
 - (a) Solutions
 - (i) Symbolic (e.g. Canada Day, statues, apologies)
 - (ii) Substantive (i.e. spending, Indigenous rights)

Descriptive



Solutions-oriented

Figure A16: Common terms in descriptive and solutions-oriented articles

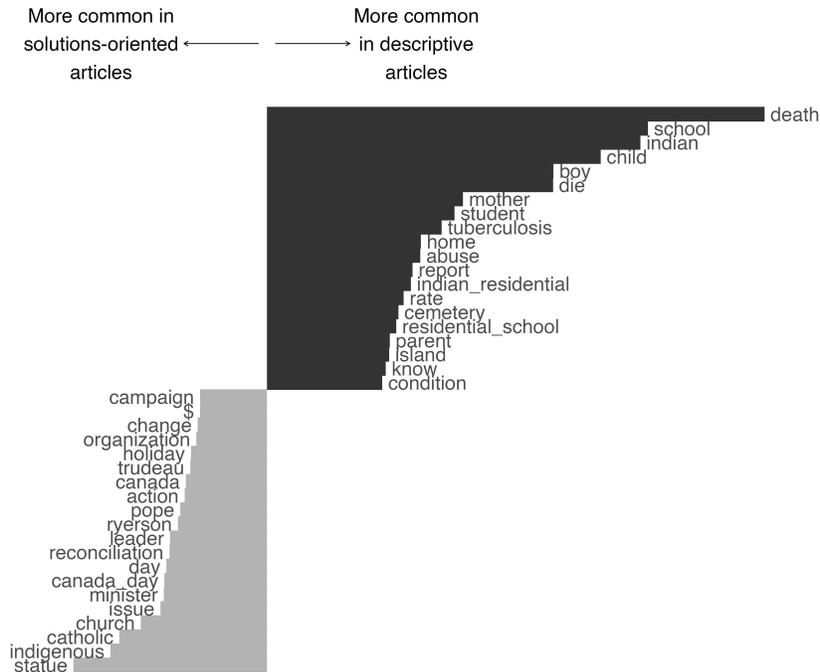


Figure A17: Relative term frequency in descriptive and solutions-oriented frames

Plot based on largest absolute χ^2 -statistics from a keyness analysis of term frequency in descriptive and solutions-oriented articles.

- (iii) Attitudinal (i.e. teaching and acknowledging the history)
- (iv) Adversarial (i.e. protests and backlash to protest methods)
- (v) Other
- (e) Responsibility attribution
- (f) Motivational calls to action

This expanded classification scheme allows me to identify the specific types of descriptive accounts and particular solutions that are more or less common.

In articles with descriptive frames, the unmarked graves and deaths are the most common focus, appearing in 55% of stories. School experiences (46%) and legacies (14%) are the next most prevalence (totals sum to more than 100% because one-quarter of all coverage contained a mix of descriptive frames). To look at how these framings change over time, Figure A18 plots the proportion of each of these sub-frames within the descriptive articles over the study period. The graves topic is especially common in the first few weeks after the initial announcement, but by the time of the second survey, it is rarely discussed in the descriptive coverage. Instead, the assimilative goals of the policy and victims' experiences in the schools have become a major focus.

In the solutions-oriented coverage, over 40% of all articles contained a focus on symbolic solutions, while only 27% and 11% discussed substantive and education-based remedies, respectively. Attributions of responsibility appeared in 19% of the solutions-oriented stories and 16% attempted to motivate the audience into supporting some form of action. Figure A19 tracks the prevalence of these sub-frames over time. Compared to Figure A18, there is less of a clear change in coverage during the study period within this set of articles. That being said, the most common solution type – symbolic – does respond in predictable ways to symbolic events like Canada Day (Week 6) and the new National Day for Truth and Reconciliation (Week 19).

Finally, a note about proportions versus counts: Figure 4 in the main text presents the proportion of articles containing each frame over the study period. Figure A20 here presents the same analysis, focusing on the total number of articles. This plot reveals that

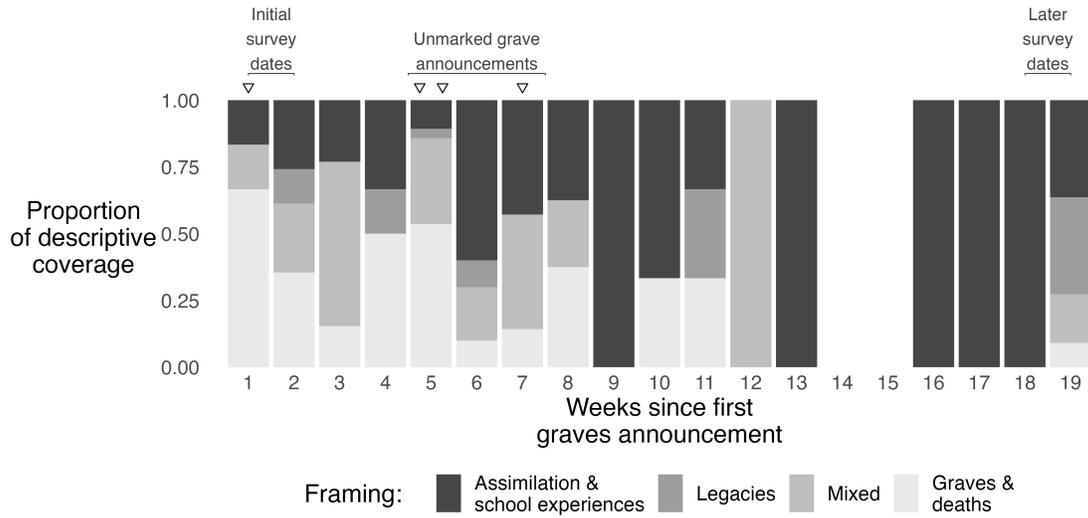


Figure A18: Changes in descriptive framings of residential schools over time

Plot presents the proportion of descriptive residential school-related articles in each week that include each topic. Note there were no descriptive articles appearing in weeks 14 and 15.

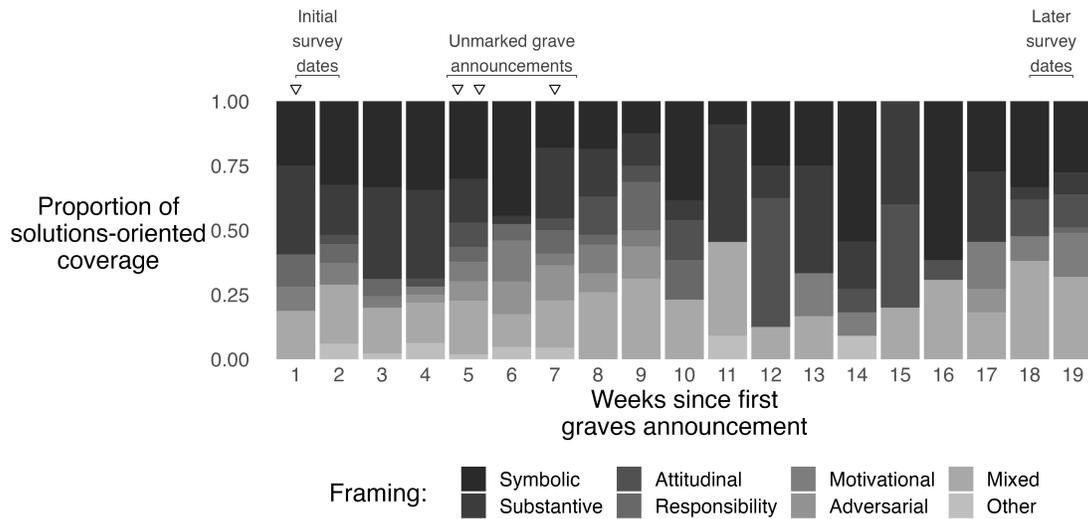


Figure A19: Changes in solutions-oriented framings of residential schools over time

Plot presents the proportion of solutions-oriented residential school-related articles in each week that include each topic.

the shift in coverage away from descriptive accounts coincided with an overall decrease in media attention to the residential schools topic. In Appendix E, I provide evidence ruling out the possibility that the reversion in attitudes toward systemic racism is driven entirely by this decline in issue salience.

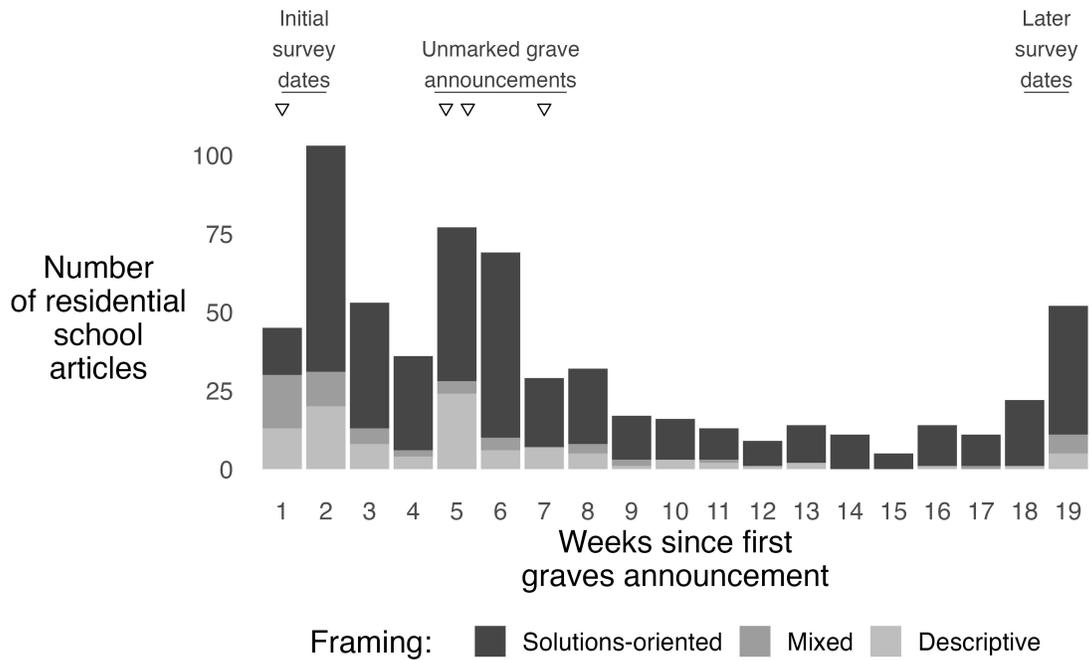


Figure A20: Prevalence of frames in residential schools articles over time
 Plot presents the number of residential school-related articles in each week that include descriptive and solutions-oriented frames.

E. CANADA’S FIRST NATIONAL DAY FOR TRUTH AND RECONCILIATION

On September 30, 2021, Canada held its first ever National Day for Truth and Reconciliation (NDTR), a federal statutory holiday that was signed into law shortly after the first unmarked graves announcement. Coincidentally, the survey dates of the Canadian Election Study overlapped with this date and, as in the main text, I compare those surveyed just before the NDTR against those surveyed just after in terms of their beliefs in systemic racism. Since the NDTR triggered a renewed attention to the history of residential schools (see Figure A22 below), this comparison seeks to descriptively capture the how an increase in the salience and social desirability of the same issues correlated with the main outcomes in a naturalistic setting. Of course, the timing of the NDTR was pre-determined and so the treatment here is not strictly exogenous (media coverage, for example, was increasing in the days leading up to the NDTR).

E.1 Balance checks

While a comparison between those surveyed just before and just after NDTR is not perfectly identified because the date of treatment exposure was not exogenous, it is nonetheless useful to check whether the two sets of respondents are similar in terms of their pre-treatment characteristics. To do so, I first compare the means on a host of presumably stable covariates in Table A11. Those in the pre- and post-NDTR samples are similar on most observable dimensions and many of the differences identified as statistically significant are small in magnitude.

Figure A21 presents a similar analysis, summarizing a model that regresses an indicator for being surveyed after the NDTR on the same set of covariates to test whether the differences persist after conditioning on other possible sample differences. The only notable result here is that a one standard-deviation increase in household income is associated with around a 1% greater likelihood of being treated. There are some other larger coefficients but they

are not distinguishable from a null of no difference. In any case, I control for all variables in Figure A21 in the analyses that follow.

Table A11: Sample characteristics by treatment status

	Average		Difference
	Pre-NDTR	Post-NDTR	
Man	0.47	0.48	0.01
Age	53.4	52.5	0.82
White	0.79	0.82	0.03*
Bachelor's degree	0.44	0.45	0.01
Household income	\$80,576	\$89,422	\$8,846*
Catholic	0.30	0.31	0.01
Other Christian	0.22	0.23	0.01
Not religious	0.38	0.37	0.01
Born in Canada	0.84	0.87	0.03*
Political interest (0 to 10)	6.45	6.47	0.02
Region: Ontario	0.35	0.32	0.03
Region: Quebec	0.29	0.30	0.01
Region: BC	0.11	0.11	0.00
Region: Atlantic	0.06	0.08	0.02*
Local Indigenous %	0.04	0.04	0.00*
French-speaker	0.25	0.28	0.03
Party ID: Bloc	0.09	0.11	0.02
Party ID: Conservative	0.25	0.28	0.03
Party ID: Liberal	0.34	0.32	0.02
Party ID: NDP	0.15	0.15	0.00
Party ID: None/Other	0.23	0.21	0.02

*p<0.05 in *t*-test for difference-in-means.

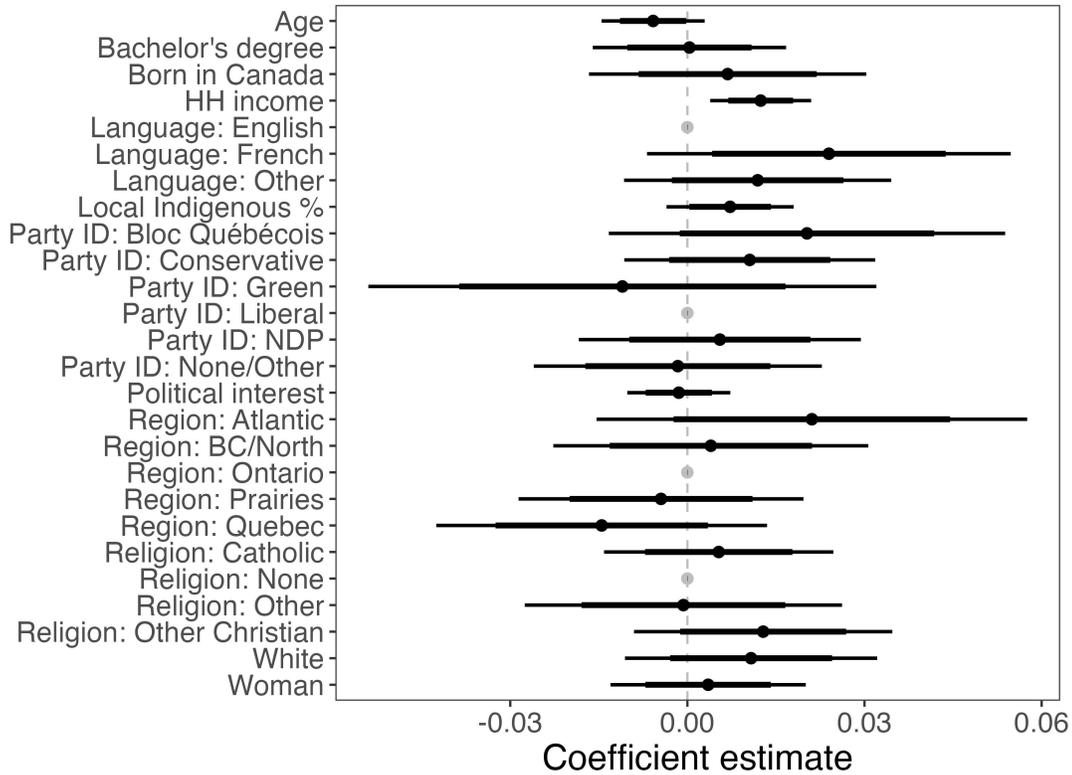


Figure A21: NDTR balance test

Plot presents coefficient estimates from an OLS model regressing treatment status on the variables listed on the y -axis. Bars indicate standard HC2 95% and Bonferroni-adjusted 95% confidence intervals. All variables are binary except for Age, Political Interest and HH income, which have been standardized such that the coefficient represents the implied effect of a one-standard deviation change. Reference categories for categorical variables are identified by grey points.

E.2 Media coverage around NDTR

To illustrate the media's renewed attention to the residential school history around the NDTR, I report in Figure A22 the number of residential school-related articles appearing by date during the 2021 Canadian Election Study. The article counts are based on the same data used in the framing analysis. While there is a moderate increase in coverage related to this topic in the day before the NDTR (September 30), there is a discontinuous jump in attention on the holiday itself. However, this media coverage is largely solutions-oriented, with descriptive and mixed framings making up only 21% of stories on the dates immediately surrounding the NDTR.

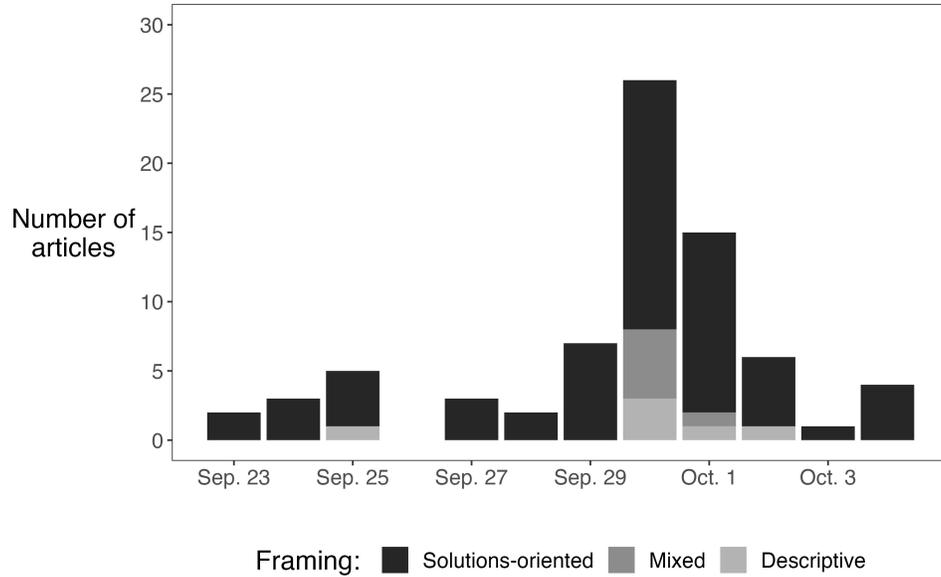


Figure A22: Media coverage of residential schools around the NDTR

Plot presents the number of residential school-related articles, by framing, that appeared on each date in the 2021 Canadian Election Study. The NDTR occurred on September 30.

E.3 Effects of NDTR on beliefs in systemic racism

To test for whether the new holiday and the media coverage that came with it had any effect on beliefs in systemic racism, I regress this attitudinal outcome on an indicator for being surveyed on or after the NDTR. Table A12 summarizes the results. Overall, there is no strong indication that attitudes improved as a result of the renewed attention to the residential school history several months after the initial unmarked graves coverage.

The results are slightly different depending on whether I employ covariate adjustment, but in the model most favourable to identifying an effect of the increase in issue salience, the renewed attention only produced around 4% of a standard deviation stronger beliefs in systemic racism. This point estimate is around 43% of the effect of the initial unmarked graves announcement and is not statistically significant, despite a sample size almost three times as large as that used in the main analysis (see Figure 1).

Overall, these results provide suggestive evidence that the reversion to baseline attitudes toward racism are not driven primarily by issues of salience or social desirability. If that were

the case, respondents surveyed just after the salience-boosting NDTR would have provided responses that more strongly endorsed the existence of systemic racism.

Table A12: NDTR and belief
in anti-Indigenous systemic racism

	Belief in systemic racism	
Surveyed on or after NDTR	-0.004 (0.034)	0.042 (0.033)
Observations	12,369	10,729
Controls	No	Yes
R ²	0.000	0.194

Coefficients are expressed in terms of pre-NDTR standard deviations. In model 2, the following covariates are included but not reported: gender, born in Canada, education, household income, party ID, political interest, religion, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. *p<0.05

F. HETEROGENEOUS RESPONSES

F.1 Partisan effect heterogeneity in quasi-experiment

One of the strongest predictors of non-Indigenous Canadians' beliefs in systemic racism is partisanship (see Figure A1). Moreover, this variable has proven an important determinant of how individuals respond to information related to racism in the American context (e.g. Chudy and Jefferson 2021; Fang and White 2022; Reny and Newman 2021). For these reasons, we might expect that supporters of different political parties may have different responses to news about unmarked graves.

To investigate whether the effects of the grave discovery differ by partisanship, I estimate conditional average treatment effects (CATEs) by interacting party identification with the post-discovery dummy in the specification from Table 1 in the main text. The results,

summarized in Figure A23, indicate that there are no substantively meaningful or statistically significant differences in treatment effects across partisan categories. There is also no evidence of a backlash effect for any subgroup: the discovery improved beliefs in systemic racism for partisans of all stripes. The only notable heterogeneity is the essentially null CATEs among supporters of minor parties and nonpartisans. In the next section, I show that this pattern is unlikely to be driven by differences in political interest.

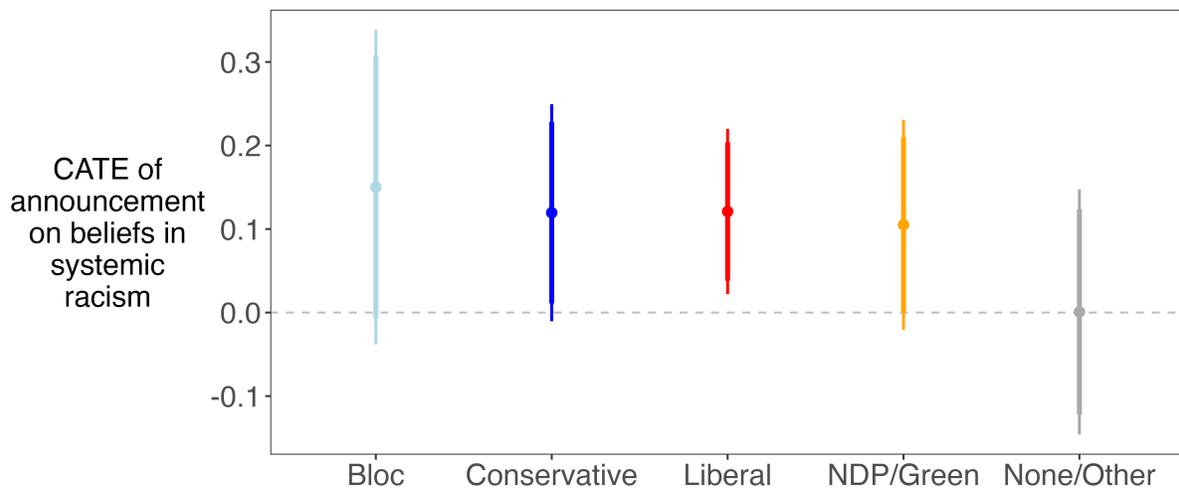


Figure A23: Effects of unmarked graves discovery by Party ID

Plot presents CATE estimates from an OLS model interacting treatment status (i.e. being surveyed after the first unmarked graves discovery) with Party ID. Coefficients are expressed in terms of pre-discovery standard deviations. Model controls for gender, born in Canada, Bachelor’s degree, log household income, party ID, turnout in 2019, Christian, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. Bars indicate 90 and 95% HC2 confidence intervals. ($n = 3,643$).

F.2 Quasi-experiment causal forest

Section F.1 investigated heterogeneity in responses to the unmarked graves news with respect to partisan identities. Yet there are numerous other subgroups that might be expected to be more or less affected by the news: white people, older people, those with less interest in the news, Catholics, immigrants and so on. I take a data-driven approach to identifying these kinds of heterogeneous treatment effects among various subgroups by training an honest causal forest (Athey and Wager 2019; Wager and Athey 2018). Causal forests are an ensemble method based on aggregating individual tree-based models that recursively partition a

“splitting” sample of the data along covariate values in order to maximize treatment effect variation within each leaf (partition of the data). Treatment effects, $\hat{\tau}_i$, are then estimated for each observation in a separate “estimation” sample of the data by assigning the mean differences in outcomes between treated and control observations within each leaf.

Figure A24 summarizes the relationships between these estimated individual-level treatment effects (on the y -axis) and several pre-treatment covariates. There are a few noteworthy patterns. First, in line with the results in Section F.1, there is little difference in treatment effects by partisanship. Second, the most notable heterogeneity can be seen with respect to the percentage of Indigenous people living in one’s local electoral district: the unmarked graves produced essentially no effect on structural racism attitudes among non-Indigenous respondents living in districts where Indigenous people made up more than 10% of the population. In fact, this variable the most important splitting criterion in the causal forest. Third, the news of the unmarked graves appears to have produced more positive effects on belief in systemic racism among those who may have had less exposure to the residential school history before 2021: immigrants, those with lower political interest, less education, and that have less trust in the media and follow the news less closely. It is also the case that men, whites and those that live in the Prairies updated their views on systemic racism less after the news broke.

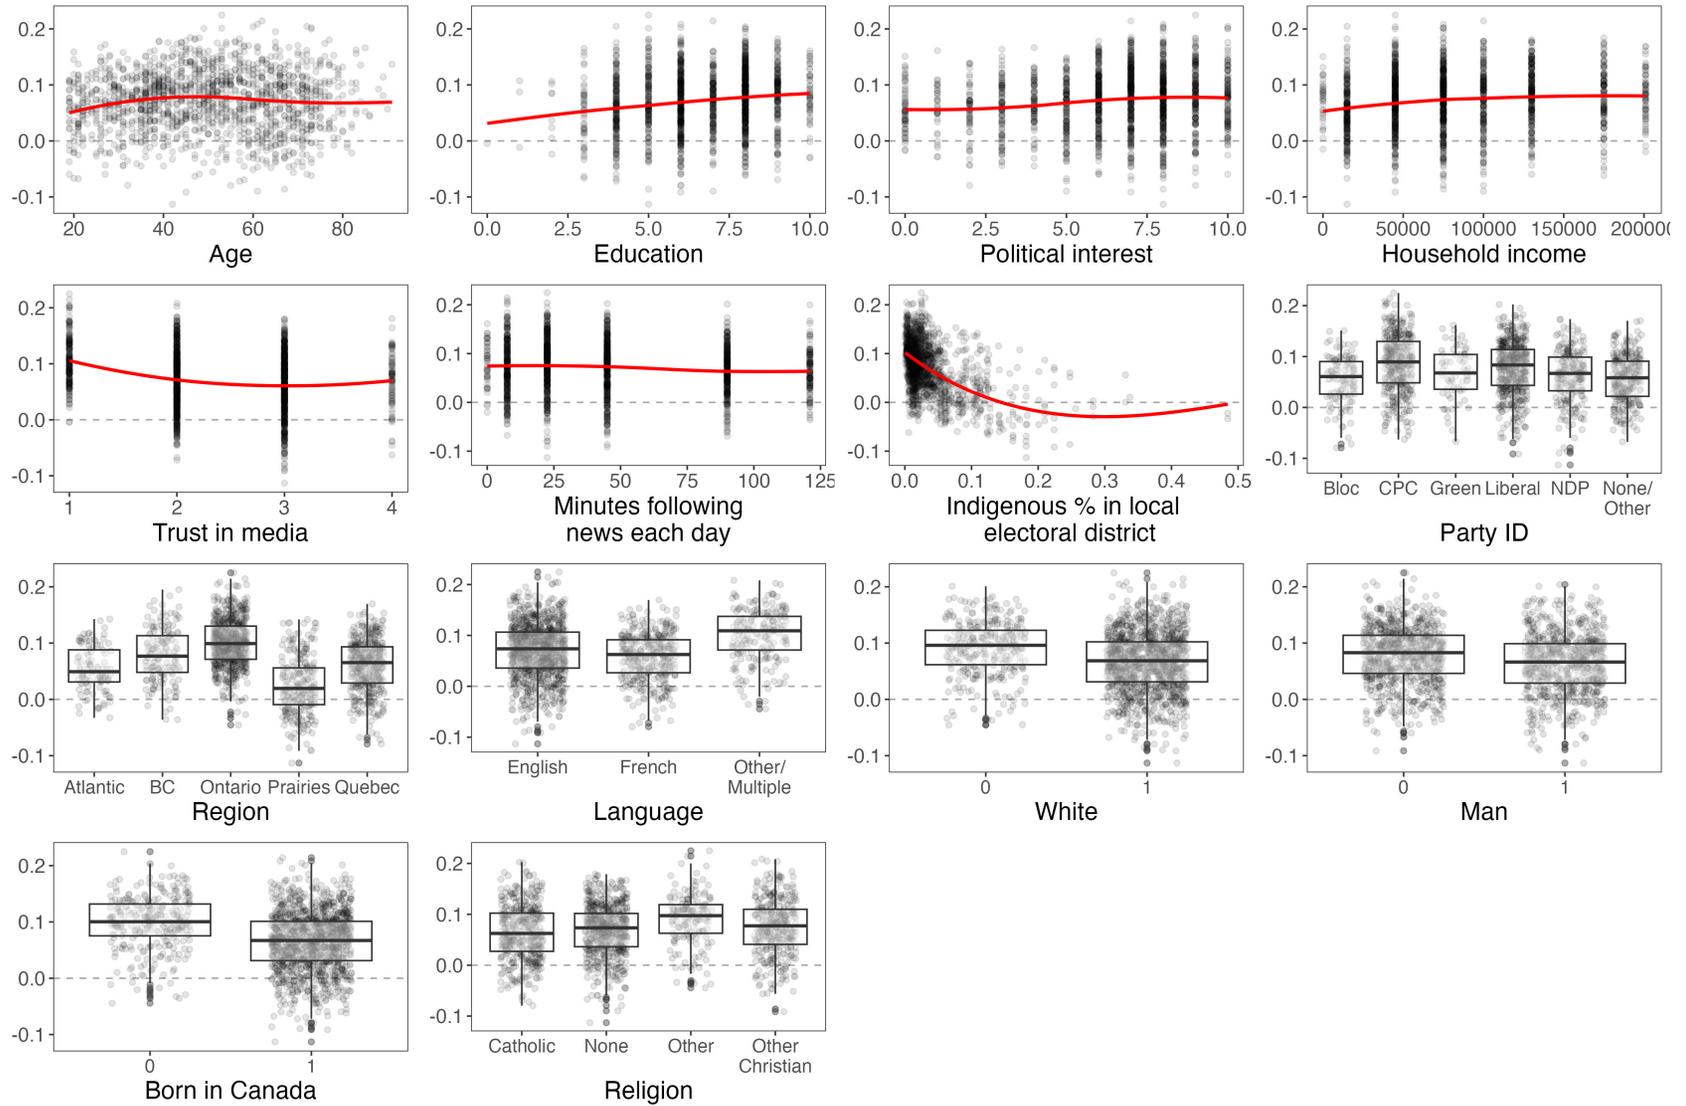


Figure A24: Causal forest estimated treatment effects and respondent characteristics

Plot summarizes the bivariate relationships between covariates and the estimated treatment effects ($\hat{\tau}$) for each observation in the test dataset. The y -axis in each plot is the estimated treatment effect.

That being said, the magnitude of these differences is quite small. Overall, there is actually little variability in units in response to the news of the unmarked graves. Figure A25 summarizes the individual-level estimated treatment effects and their variability. Most $\hat{\tau}$ values fall within 0 to 0.2 s.d. and there is a fair degree of uncertainty around these estimates. Perhaps most importantly, 90% of all observations were estimated to have a positive treatment effect, meaning there is little evidence of backlash among particular subgroups.

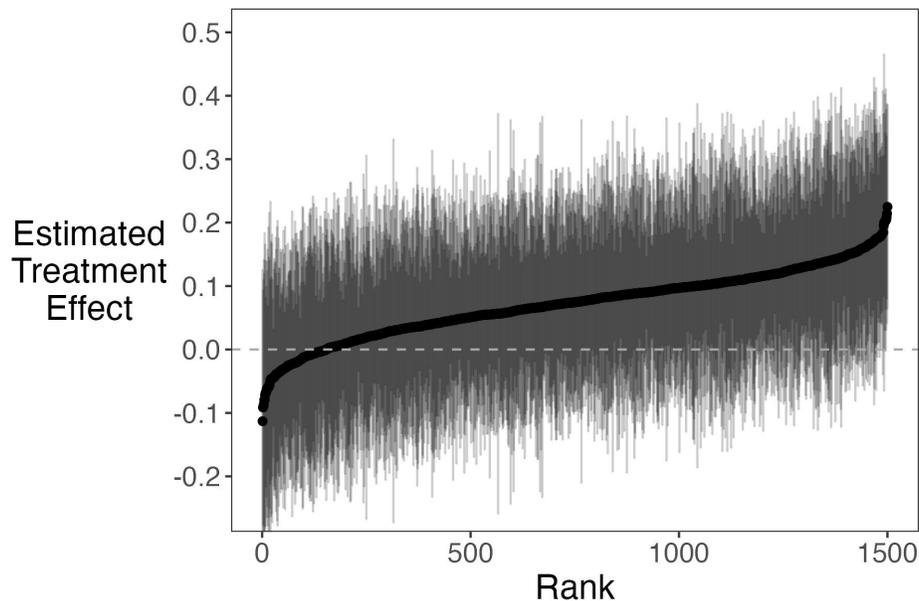


Figure A25: Causal forest estimated treatment effects

Plot presents the estimated treatment effect ($\hat{\tau}$) for each observation in the test dataset, ranked by their $\hat{\tau}$ values.

Finally, the bivariate plots earlier ignore correlations among covariates and do not capture estimation uncertainty. An alternative summary CATE measure is the best linear projection (BLP), a doubly robust estimate of the following linear model:

$$\tau(\mathbf{X}_i) = \alpha + \mathbf{X}_i\beta$$

where $\tau(\mathbf{X}_i)$ is the CATE and \mathbf{X}_i is a vector of covariates. Of course, this modelling approach assumes linearity between the CATEs and covariates, which may not be strictly true given the relationships in Figure A24, but it does offer a useful starting point for characterizing

effect heterogeneity.

Figure A26 summarizes the coefficients from the BLP. Most of the patterns in the bivariate plots hold up in this analysis: while there is not a significant amount of effect heterogeneity, some subgroups do exhibit slightly different responses to the treatment. For example, women were more affected by the news, while those have less trust in the media, lower political interest and have a larger Indigenous population in their area updated their beliefs less positively about the existence systemic racism.

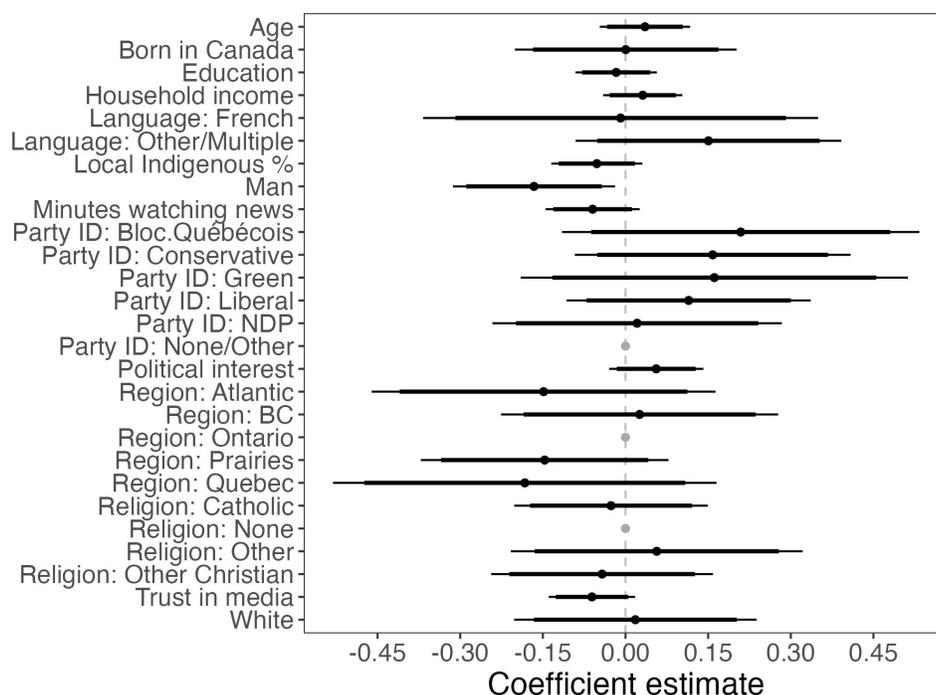


Figure A26: Best linear projection for causal forest estimated treatment effects

Plot presents the coefficient estimates from a best linear projection of the CATEs on covariates listed in the y -axis. All covariates are categorical except age, education, household income minutes watching news and trust in media, which have all be standardized to represent implied effects in terms of standard deviation changes.

F.3 Local demographic CATEs in quasi-experiment

The causal forest described in Section F.2 identified the size of the Indigenous population in a respondent's electoral district as an important determinant of the effects of the unmarked grave discoveries. But the relationship did not appear entirely linear, making the best linear

projection in that section potentially misleading. In Figure A27, I instead summarize the CATEs from an OLS model interacting the treatment indicator with a dichotomized version of the local Indigenous population variable. The point estimates suggest that the news of the unmarked graves increased beliefs in structural racism among respondents living in areas where less than 10% of the population is Indigenous, but not in areas above that threshold. However, because only 10% of the sample actually lives in areas with sizable Indigenous communities, there is a large degree of uncertainty around the estimate for those in electoral districts with an Indigenous population share above 10% and we cannot confidently reject the possibility that the effect is the same for respondents in both areas. The possible effect heterogeneity here would be a promising line of inquiry for future research.

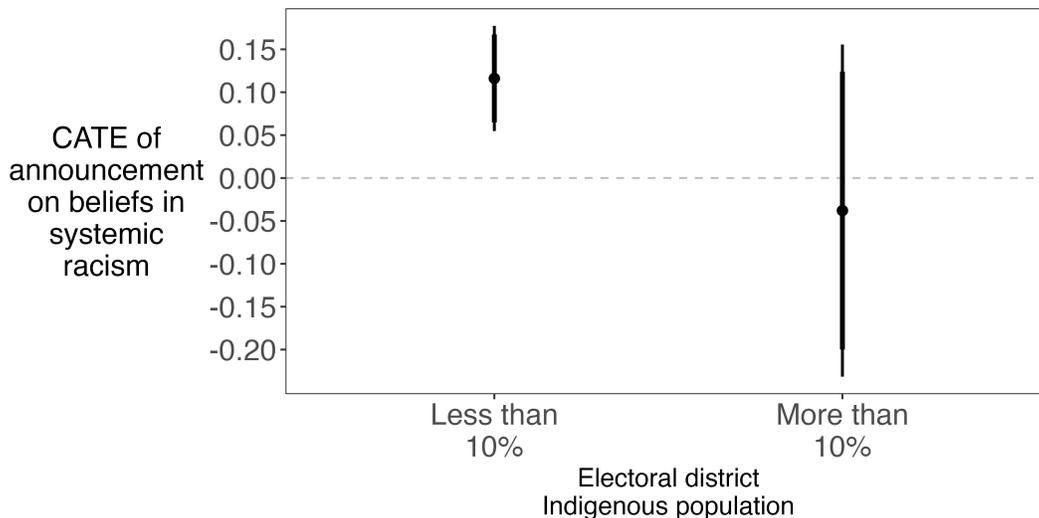


Figure A27: Effects of unmarked graves discovery by local Indigenous population

Plot presents CATE estimates from an OLS model interacting treatment status (i.e. being surveyed after the first unmarked graves discovery) with a dichotomized indicator of the size of the Indigenous population in a respondent's electoral district. Coefficients are expressed in terms of pre-discovery standard deviations. Model controls for gender, born in Canada, Bachelor's degree, log household income, party ID, turnout in 2019, Christian, language, white, electoral district Indigenous percentage, province, and birth-decade fixed effects. Bars indicate 90 and 95% HC2 confidence intervals. ($n = 3,643$).

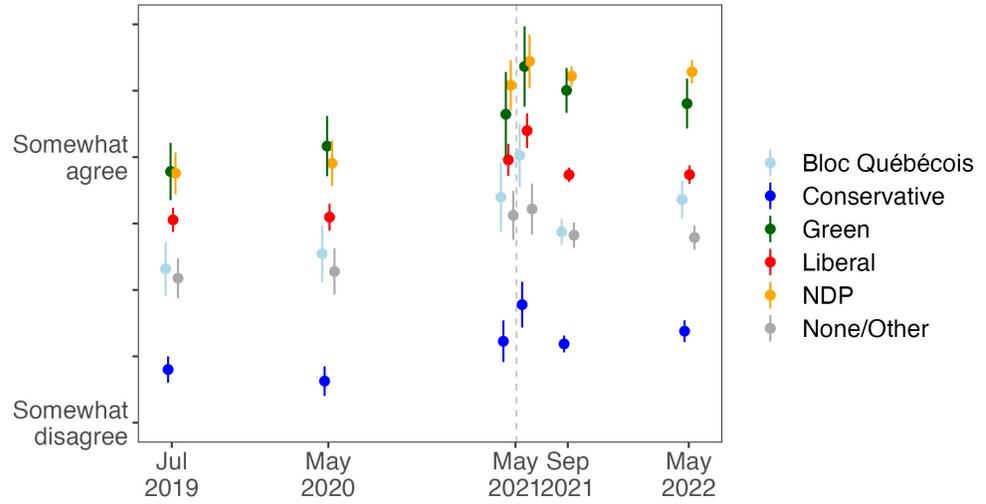
F.4 Heterogeneity in attitudinal persistence

In Sections F.1 and F.2, I show that there is minimal heterogeneity in responses to the initial news of the unmarked graves. In this section I show that there are also no major

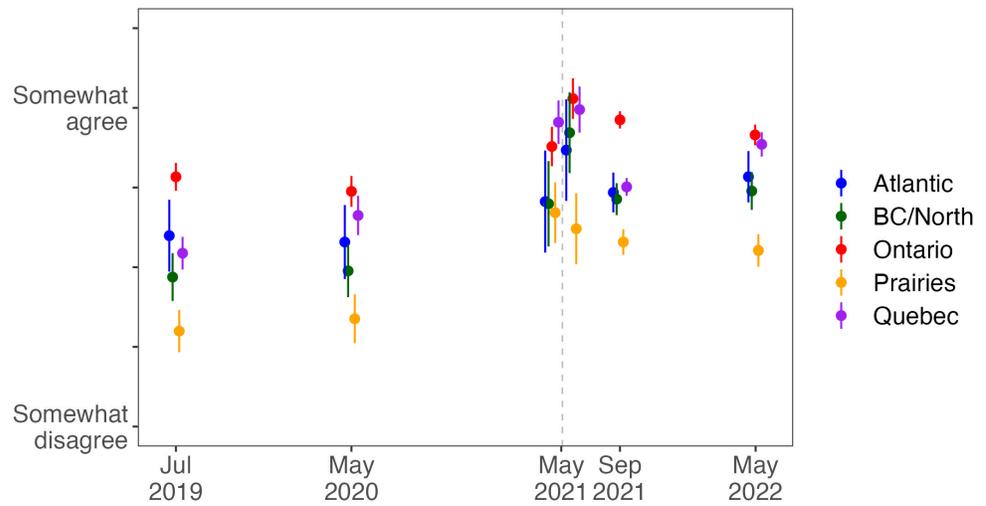
differences in patterns of attitudinal persistence across relevant subgroups. Figure A28 plots the average agreement with the two systemic racism items over time separately across values of several different pre-treatment variables. Recall that in Figure 3 from the main text, after an increase in beliefs in the existence of systemic racism after the initial unmarked graves news in May 2021, levels of agreement with the two items returned to baseline and possibly even worsened in a September 2021 survey and did not change in a final survey in May 2022.

Overall, this pattern is generally the case for all subgroups in Figure A28 as well. There are some differences in the magnitude of attitudinal decay or persistence, but almost every group returned to baseline or worse in the September 2021 survey. One exception is the group of atheists and agnostics, but given the number of comparisons being made here, it is not surprising that at least one subgroup would display an unexpected pattern and so it is difficult to know if this result is just noise. The larger takeaway is that on average, non-Indigenous Canadians with various background characteristics responded in the same way to the initial news and then reverted to prior beliefs in similar ways.

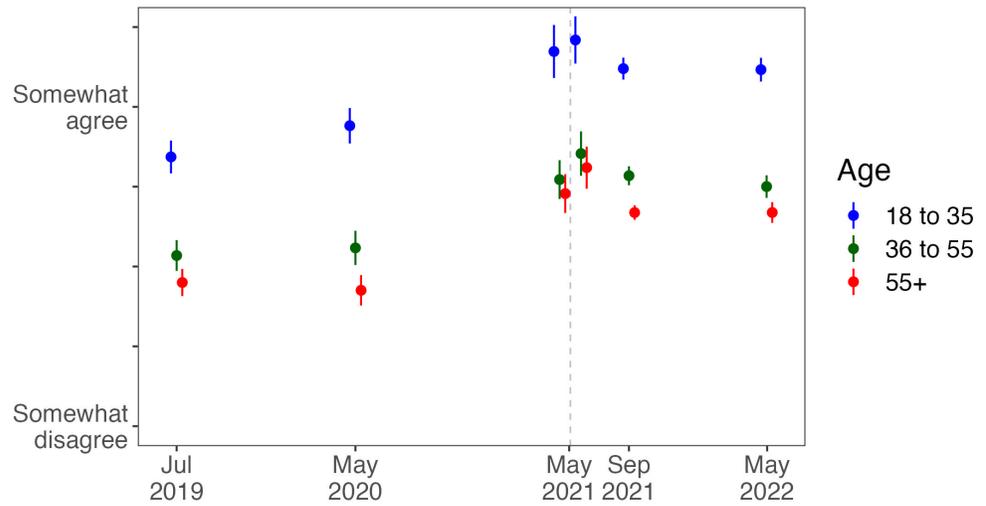
Agreement with existence of structural racism



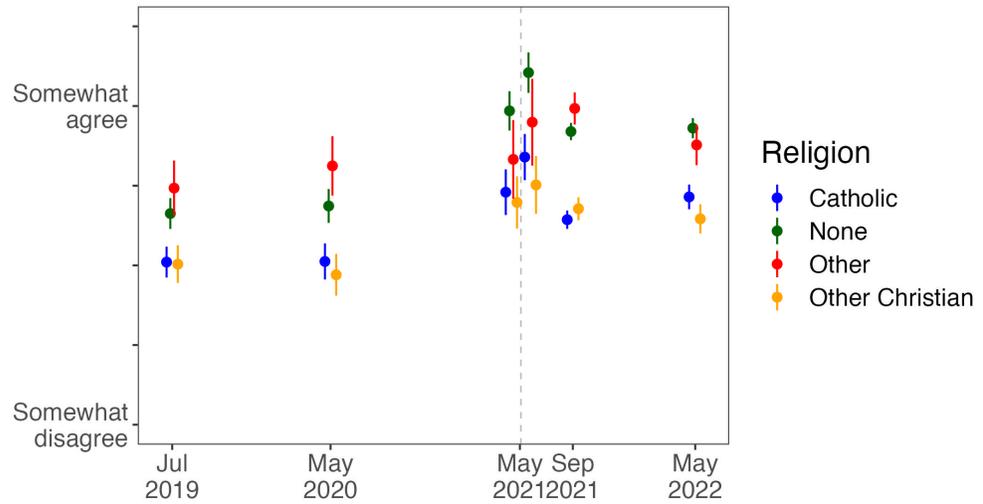
Agreement with existence of structural racism



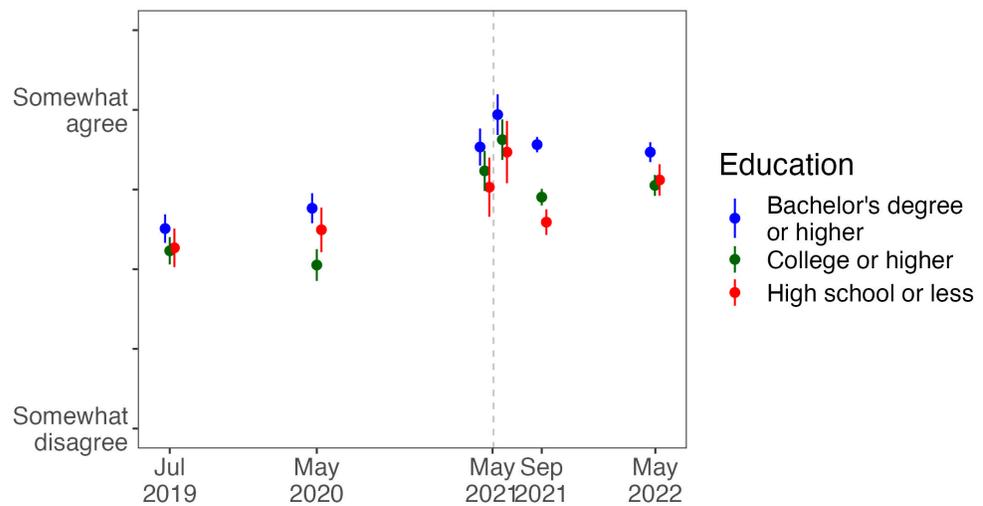
Agreement with existence of structural racism



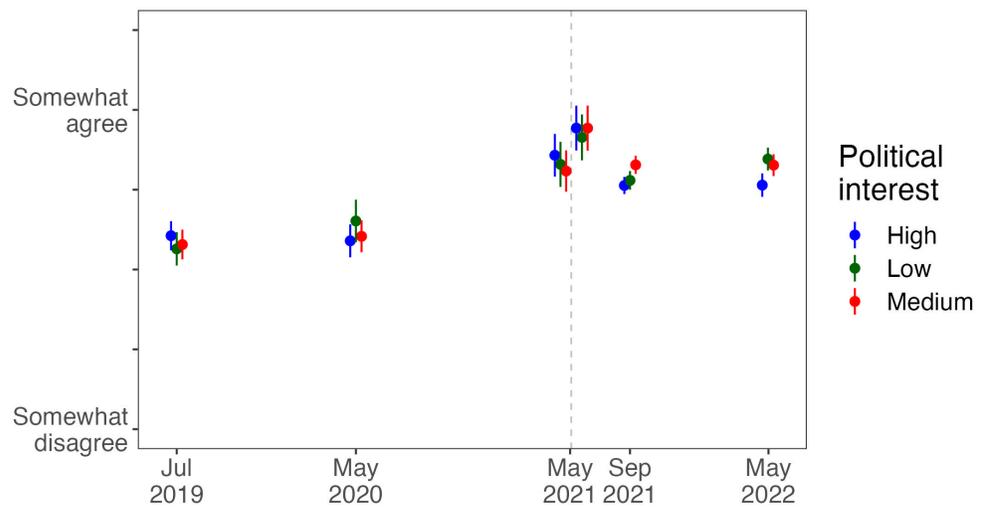
Agreement with existence of structural racism



Agreement with existence of structural racism



Agreement with existence of structural racism



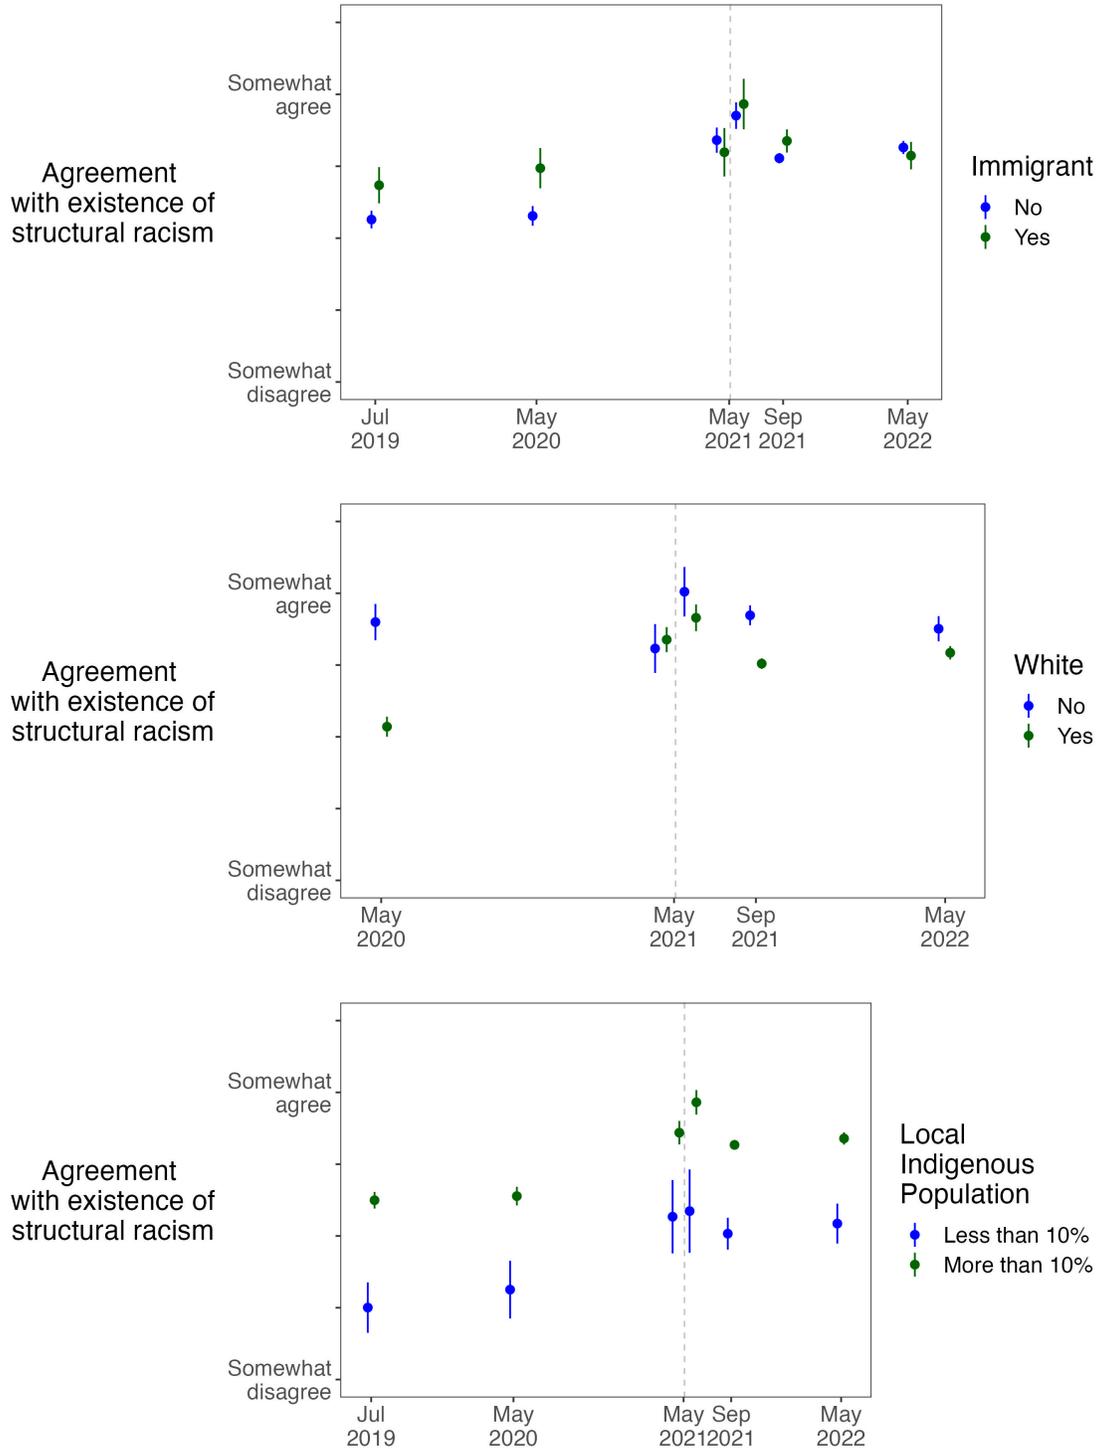


Figure A28: Agreement with systemic racism by covariates, 2019 to 2022

Plot presents average and 95% confidence intervals for the mean of a respondent-level average of the two items measuring beliefs in systemic racism (see Appendix Table A6 for each item plotted separately). Higher values indicate greater belief in the existence of systemic racism. In all but the September 2021 survey, this question was asked on a four-point Likert scale. For that specific survey, a five-point scale was used, but responses have been rescaled to match the four-point scale (see Appendix Figure A7 for individual response level prevalence over time).

REFERENCES

- Athey, Susan, and Stefan Wager. 2019. Estimating treatment effects with causal forests: An application. Publisher: University of Pennsylvania Press, *Observational Studies* 5 (2): 37–51.
- Benford, Robert D., and David A. Snow. 2000. Framing Processes and Social Movements: An Overview and Assessment. *Annual Review of Sociology* 26 (1): 611–639.
- Chudy, Jennifer, and Hakeem Jefferson. 2021. Opinion — Support for Black Lives Matter Surged Last Year. Did It Last? *The New York Times* (May).
- Fang, Albert H., and Steven White. 2022. Historical information and beliefs about racial inequality. *Politics, Groups, and Identities*.
- Grimmer, Justin, and Brandon M. Stewart. 2013. Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political analysis* 21 (3): 267–297.
- Reny, Tyler T., and Benjamin J. Newman. 2021. The opinion-mobilizing effect of social protest against police violence: Evidence from the 2020 George Floyd protests. *American Political Science Review* 115 (4): 1499–1507.
- Schofield, Alexandra, and David Mimno. 2016. Comparing apples to apple: The effects of stemmers on topic models. *Transactions of the Association for Computational Linguistics* 4:287–300.
- Wager, Stefan, and Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* 113 (523): 1228–1242.