

DIGITAL DEMOCRACY PROJECT

Research Memo #5 Fact-Checking, Blackface and the Media

OCTOBER 2019

Produced by

Taylor Owen, Project Director

Peter Loewen, Head of Survey Analysis Team

Derek Ruths, Head of Online Data Analysis Team

Aengus Bridgman

Robert Gorwa

Stephanie MacLellan

Eric Merkley

Andrew Potter

Beata Skazinetsky

Oleg Zhilin



DIGITAL DEMOCRACY PROJECT

The Digital Democracy Project is a joint initiative led by the Ottawa-based Public Policy Forum and the Max Bell School of Public Policy at McGill University.

The project will study the media ecosystem in the run-up to and during Canada's October 2019 federal election by monitoring digital and social media and by conducting both regular national surveys and a study of a metered sample of online consumption. The project will communicate its preliminary research findings publicly on a regular basis from August to October 2019, and will work with journalists to analyze the spread and impact of misinformation. The study will culminate in a final report to be published by March 2020. Both the project's preliminary findings and final report will be publicly available.

The project director is Taylor Owen, Associate Professor and Beaverbrook Chair in Media, Ethics and Communications in the Max Bell School of Public Policy at McGill University. The online data analysis team is led by Derek Ruths, Associate Professor in the School of Computer Science at McGill University, and the survey analysis team is led by Peter Loewen, Professor of Political Science at the University of Toronto.

The project is funded by The Rossy Foundation, the McConnell Foundation, and the Luminare Group and with support from the Mozilla. The project is also participating in the Digital Elections Research Challenge, a collaborative research project led by Taylor Owen and Elizabeth Dubois, Assistant Professor at the University of Ottawa, and funded by a grant from Heritage Canada. The DDP will be sharing survey and online data with the 18 research projects funded through this collaboration and will highlight select findings from these projects in our regular briefings.

For enquiries, please contact [Stephanie MacLellan](#).

TABLE OF CONTENTS

Executive Summary 2

Methodology 5

General Findings: Support for Fact-Checking in Canada 6

Findings: Effectiveness of Different Fact-Checking Sources 10

Findings: Trust and Internal Efficacy 13

Case Study: Trudeau and Blackface 16

Findings: Emergence and Decline of Blackface Hashtags 16

Findings: Partisan Differences and Isolation 22

Findings: Opportunism and Manipulation 25

Appendix 27

EXECUTIVE SUMMARY

If there is a general theme emerging from these Digital Democracy Project reports, it is that much of the received wisdom about the current state of the media and democracy is, if not false, at least poorly supported by the available evidence. For example, contrary to popular belief, Canadians do not live in media-driven partisan echo chambers, and their trust in mainstream media remains quite high. And while there is some evidence of political polarization, that appears to be driven largely by the parties themselves and not by the media.

The good news, then, is that there appears to be a great deal of room for the media to play a valuable and constructive role in sustaining good-faith democratic engagement. This week's report builds on these themes in two ways.

First, we have new survey results that suggest that there is a great deal of support among Canadians for fact-checking, and that it can be effective. Yet while fact-checking is a practice that has long been considered part of the bread and butter of accountability journalism, there is little evidence that journalist-based fact checking is more effective than any other means.

Second, a data-driven analysis of Justin Trudeau's blackface controversy shows that the initial flood of online activity was marked by widespread sharing of solid journalism. As the initial interest in the story waned, further activity consisted by and large of conservatives talking among themselves, with little evidence of the conversation being influenced by inauthentic actors.

KEY ISSUES: FACT-CHECKING

- 1** Canadians are in broad agreement that fact-checking is valuable in politics. This echoes a key finding from [Research Memo 1](#), which found high levels of trust in traditional media to deliver objective and accurate information about the election.
- 2** Our finding from [Research Memo 2](#), that fact-checking works in correcting misinformed voters, was further confirmed in this report. However, this report found that journalists are no more convincing than politicians or unaffiliated Twitter users when it comes to fact-checks.
- 3** Fact-checks may help people to feel more confident in their ability to understand and participate in politics, but this is more likely to be the case for those who already understand the issues. We also find no evidence that fact-checking politicians reduces trust in politicians as a source of political information.

KEY ISSUES: TRUDEAU AND BLACKFACE

- 1 **The general public’s discussion of the blackface story on social media dropped dramatically after three days. Tweets on the topic from journalists and election candidates had a similar decline.**
- 2 **Discussion on blackface-related hashtags is dominated by Conservative partisans; features links to generally informative pieces of journalism about the controversy; and lacks evidence of disproportionate inauthentic activity.**

Fact-checking as a journalistic practice has grown in prominence in the past few years, especially since the election of an American president who habitually bends or outright ignores the truth. It’s estimated the number of fact-checking outlets around the world has [quadrupled since 2014](#), with a 26% increase in the past year. But as fact-checking gains public profile, so do arguments about whether it is effective in informing the public or changing the minds of people who are inclined to reject information that contradicts their beliefs.

Regardless of the answers to those questions, the Canadians we surveyed were broadly in favour of fact-checking. Fifty-three percent of respondents said they supported the practice, while only 14% expressed any level of opposition. An overwhelming 73% of respondents indicated they wanted more of the practice in journalism, while only 3% indicated they wanted less of it. And there were only minor differences in support between left- and right-leaning partisans, or between those who were more or less exposed to traditional or social media.

Canadians clearly want to see more fact-checking, but how effective is it when it comes to informing the public? The Digital Democracy Project has been using survey research to study how Canadians respond when they are provided with facts related to policy issues. In [Research Memo 2](#), we found that survey respondents who received information about Canada’s progress on its Paris Accord commitments were more likely to answer a related question correctly, regardless of their partisan leanings.

This report builds on those earlier findings by tracking Canadians’ responses to fact-checking scenarios. We provided statements from a fictional MP that included false information about Canada’s greenhouse gas emissions and violent crime rate. Respondents then received a fact-check from either a journalist, another politician or an unaffiliated Twitter user (all fictional), or no fact checks at all. We found that the fact checks made a difference: For the crime rate question, 55% of respondents who did not receive the fact-check answered incorrectly, compared to 43% who did—a difference of 12 percentage points. For the question about emissions, the share of respondents answering incorrectly was almost 18 points lower among those who received the fact-check than among those who did not.

But perhaps surprisingly, **the source of the fact-check made less of a difference.** In [Research Memo 1](#), respondents reported trusting the traditional news media as a source of political information at a rate of 5.6 out

of 10, surpassing their trust in politicians (4.8) and social media (3.3). However, in this survey, there was no significant difference in responses to fact-checks from the journalist, the politician or the random Twitter user. If journalists want to convince their audiences that they are the most reliable source for verifying information, they may have more work to do.

Finally, we used a second fact-checking scenario to evaluate the effect on respondents' *internal efficacy*—or their belief that they can understand and participate in politics. Respondents were again given inaccurate claims on four different topics from a fictional politician, and anywhere from zero to four fact-checks on those claims from a fictional journalist. Respondents who knew the correct information about the politicians' claims at the outset of the experiment saw their efficacy improve as they read more fact-checks. By contrast, those who did not know the information ahead of time saw no significant change in their efficacy. This suggests that fact-checking may actually increase the inequality in internal efficacy between those with high and low levels of policy knowledge.

SOCIAL MEDIA FINDINGS

Our social media case study focused on the controversy around the images of a younger Prime Minister Justin Trudeau in blackface and brownface that surfaced one week into the election campaign. When a political story emerges as suddenly as this one did, it allows us to observe how social media behaviour changes in response.

Indeed, there was a dramatic spike in tweets and Facebook posts about blackface and brownface in the 24 hours after the story broke, but the decline was almost as dramatic. At its peak on Sept. 19, there were more than 13,400 tweets an hour on the topic; three days later, there were no more than 2,614 tweets an hour. Journalists and politicians—particularly NDP and People's Party of Canada candidates—continued to tweet about the issue for the rest of the week, albeit in much smaller numbers. Our previous reports have shown that political insiders and journalists are not always having the same conversations as the rest of the public on social media. In this case, the dwindling conversations in all three spheres seemed to reflect each other.

We also found that a substantial amount of Twitter activity around related hashtags such as #blackface or #brownface came from Conservative partisans. We identified Twitter users' likely partisan affiliations by looking at the number of candidates they followed. About 11% of Conservative partisans tweeted or retweeted blackface-related hashtags, the highest for any party. By comparison, 10% used right-wing hashtags (e.g. #trudeaumustgo, #scheer4pm) and 16% used general election hashtags (e.g. #cdnpoli, #elxn43) over the same time period. Not surprisingly, Liberal partisans continued to use general election hashtags far more frequently than blackface-related hashtags, but so did the NDP, Bloc Québécois and Greens. PPC partisans continued to promote right-wing hashtags in greater numbers than blackface hashtags.

A similar picture emerged when we looked at co-occurrence of hashtags: how often certain election-related hashtags coincide with each other in the same tweet. Hashtags can allow Twitter users to go beyond their usual

online networks by seeing what anybody else is posting on that hashtag. The #blackface hashtag was used by more Conservative partisans than anyone else, and on top of that, most of the hashtags that interacted with #blackface were also dominated by Conservative and PPC partisans. This suggests that when it comes to social media conversations about the controversy, right-leaning partisans are mainly talking among themselves.

METHODOLOGY

Our survey data team conducted an online panel survey of 1,594 Canadian citizens 18 years and older using the online sample provider Qualtrics. The sample was gathered from Sept. 19-24. Data was weighted within each region of Canada by gender and age to ensure it adequately represented the Canadian public. Survey respondents were asked questions related to basic demographics, as well as their partisan, ideological and issue preferences. We present 90% confidence intervals for each of our figures below. Partisan sub-groups are restricted to the Conservatives, Liberals and NDP for sample size considerations.

For our online data research, we continue to track data from across the digital ecosystem. This week we rely on approximately 3 million tweets collected since Sept. 17. We gathered a list of all the Twitter followers of all declared major party candidates (as of Sept. 28). We identified active partisan consumers in the election, which we defined as individuals who follow more than five rank-and-file candidates from major parties (we exclude the top 5% most-followed individuals from each party to account for the very high number of followers of some party leaders, where a follow does not meaningfully indicate partisanship). We labelled each of these approximately 55,000 Twitter users as likely partisans based on the party that includes most of the candidates they followed. We also used text and links from candidates, news outlets, and other public groups and pages on Facebook.

More details about our methodology can be found in the Appendix.

GENERAL FINDINGS: SUPPORT FOR FACT-CHECKING IN CANADA

We asked our respondents a pair of questions to gauge their support for fact-checking in journalism. Our respondents are overwhelmingly in favour of the practice. Fifty-three percent indicated they favoured the practice, while only 14% expressed any level of opposition. An overwhelming 73% indicated they wanted more fact-checking in journalism, while only 3% wanted less of it. Canadians are in broad agreement that fact-checking is valuable in politics.

Our survey allows us to explore factors that might drive support or opposition to fact-checking. We averaged our two measures of fact-checking support to construct a scale that runs from 0-1, where 1 is the most supportive of fact-checking. Media trust appears to be linked to support for fact-checking, which is not surprising because journalists are the main source of fact-checks. We measured trust in the media in this wave with an index of four questions scaled from 0-1, where 1 is the most trusting of the news media.¹ Respondents who were more trusting of the media offered more support for fact-checking (0.79) than those who were less trusting (0.72). There also appear to be some modest partisan differences in support for fact-checking. Left-leaning partisans are more supportive of fact-checking (0.77) than right-leaning partisans (0.72), which might be in part a function of differing levels of trust in the media—left-leaning partisans are modestly more trusting of the news media (0.49) than right-leaning partisans (0.43).

¹ 1) I trust the mass media – such as newspapers, TV, and radio – to report the news fully, accurately and fairly; 2) The news posted online by news organizations is accurate; 3) I have confidence in the people running the press; 4) When dealing with political and social issues, news organizations deal fairly with all sides (strongly agree to strongly disagree; 5-point)

FIGURE 1. Favourability of fact-checking (left panel); desire for more or less fact-checking by journalists (right panel)

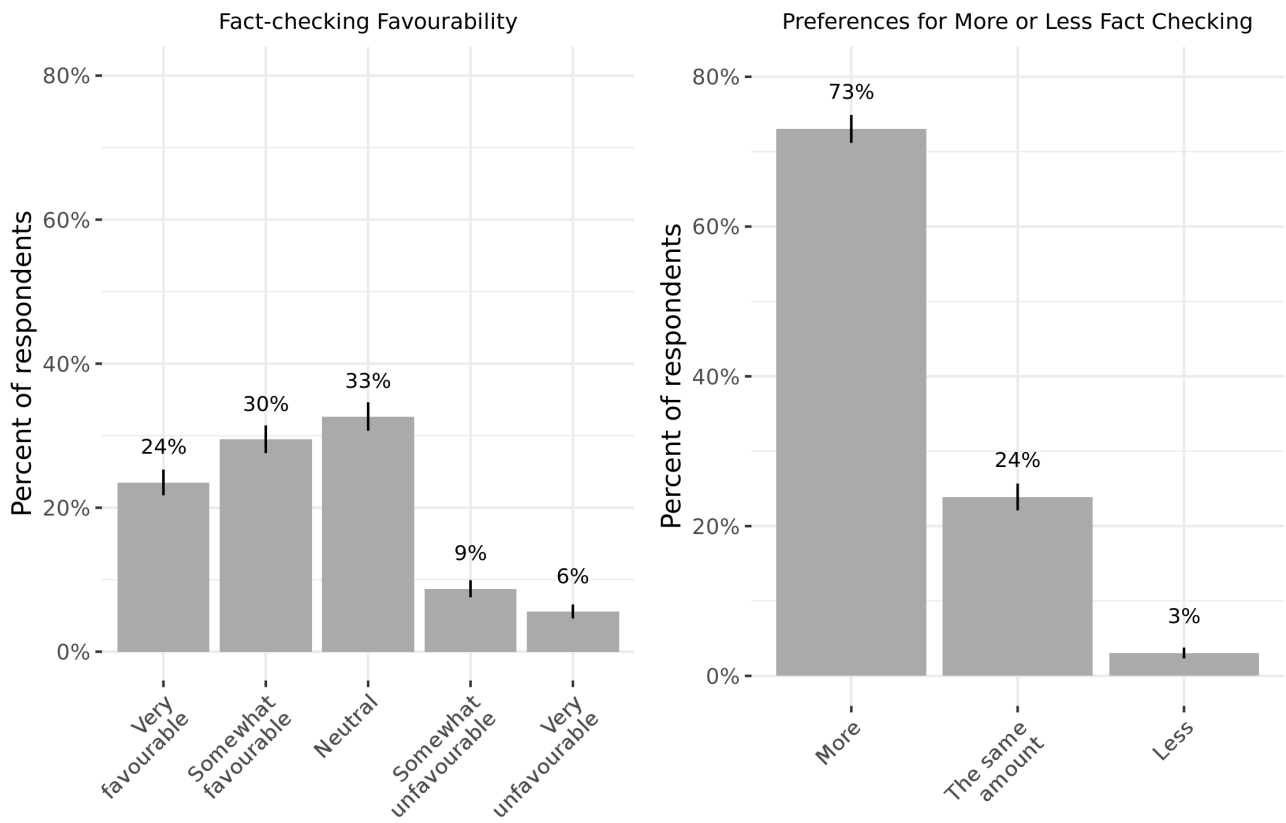


TABLE 1. Determinants of Support for Fact-checking

	Support
Left	0.77
Right	0.72
Difference	0.05
High Media Trust	0.79
Low Media Trust	0.72
Difference	0.07
High Political Knowledge	0.78
Low Political Knowledge	0.72
Difference	0.06
High Political Interest	0.77
Low Political Interest	0.72
Difference	0.05
Strong partisans	0.76
Weak partisans	0.77
Difference	-0.01
High Traditional News Exposure	0.76
Low Traditional News Exposure	0.72
Difference	0.04
High Social Media Exposure	0.71
Low Social Media Exposure	0.75
Difference	-0.04
Partisan Media Exposure	0.71
No Partisan Media Exposure	0.76
Difference	-0.05

As Table 1 shows, fact-checking is also comparatively more supported among those with higher levels of political knowledge, and those with higher levels of political interest.² Interestingly, there are only trivial differences in support for fact-checking between strong and weak party supporters. Traditional news media exposure is associated with higher levels of support for fact-checking, while exposure to social media and to media sources that draw a highly partisan audience are both associated with lower support for fact-checking.³ However, it is important to note that these differences are modest. Across all groups of respondents, fact-checking commands overwhelming support.

Finding 1: Fact-checking enjoys widespread support in the Canadian public, and this support reaches across different segments in society. There is some evidence that right-leaning partisans are slightly less supportive of fact-checking than left-leaning partisans, as are those with lower levels of trust in the media. Exposure to media sources favoured by a partisan audience, however, appears to be strongly associated with lower support for fact-checking.

² Respondents were asked a series of fact-based questions about politics, including the unemployment rate, the name of the Secretary General of the United Nations, and the relative ideological placement of the political parties. Respondents in the bottom third of the distribution were classified as low in political knowledge, while respondents in the top third were labelled as high in knowledge. Respondents rated their general interest in politics on a 0-10 scale. The distinction between high and low interest was made by splitting respondents at the median of the scale.

³ We are tracking a self-reported exposure to a number of different traditional news outlets and social media applications, which is described more fully in our first report. We classify those who have high exposure as being in the top third of the distribution and those who have low exposure as being in the bottom third. We identify partisan-congenial media sources as those that are selectively followed and shared by partisans on Twitter, as in [Research Memo 3](#).

FINDINGS: EFFECTIVENESS OF DIFFERENT FACT-CHECKING SOURCES

In [Research Memo 1](#), we asked our survey respondents who they trusted to deliver objective and accurate information about the upcoming federal election. Respondents trusted the traditional news media (5.6) at levels that approached their family and friends (6.0) on a scale of 0-10, with 10 being the highest. They trusted politicians (4.8) and the information they received on social media (3.3) far less. Among the news outlets we measured, CBC was the most trusted (6.2). These findings suggest that Canadians are attentive to the *sources* of information they encounter in political discourse. They trust journalists from established outlets such as the CBC and may be more willing to accept information provided by them, as opposed to politicians or people they encounter on social media.

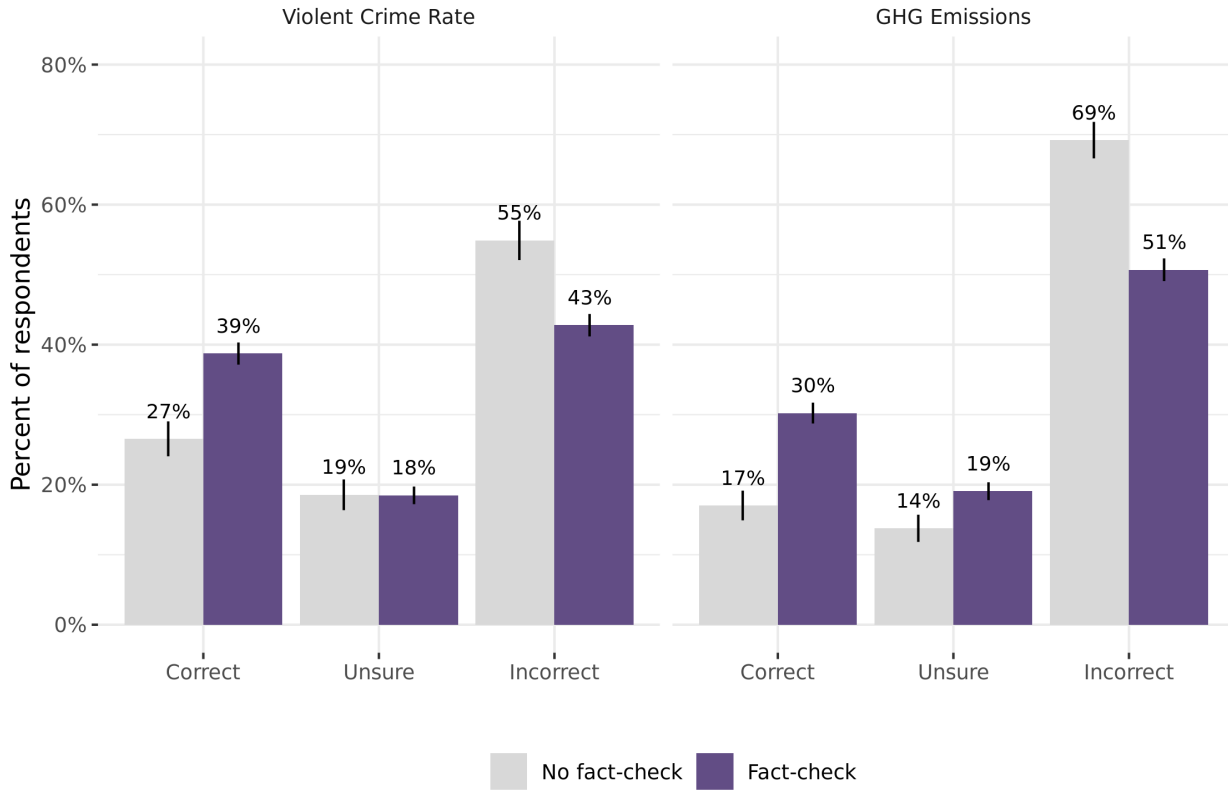
In [Research Memo 2](#), we examined the effectiveness of a fact-check on whether Canada is on track to meet its Paris Climate Accord commitments (answer: no). We found respondents who received the correct information were more likely to answer a related question correctly, regardless of their partisan leanings.

In this report, we examine whether the *source* of the fact-check has any influence on its effectiveness. In our survey run from Aug. 28 to Sept. 5, we randomly assigned respondents to receive fact-checks on two issues: the trend in Canada's greenhouse gas emissions since 2015 (answer: lower), and the current violent crime rate compared to the past decade (answer: lower). In both fact-check scenarios, the source of the misinformation was a fictional member of Parliament.

All respondents were provided with both incorrect facts, and they were randomly assigned to receive a fact-check from one of three sources—a CBC journalist, an MP or an unaffiliated person on Twitter—or to receive no fact checks at all. (All actors in these fact-checking scenarios were fictional.) After receiving the information, respondents were asked to rate the accuracy of the politician's original claim, or to say they were unsure.

Ideally, fact-checks should work at correcting misperceptions, but only when the source of that fact-check is credible. Based on the results of our first survey, we would expect the fact-check from the journalist to be most effective. Fact-checks from social media users and politicians should be less effective, or perhaps not effective at all given the lack of credibility our previous respondents assigned to these sources.

FIGURE 2. Effectiveness of fact-checks on respondent perceptions toward the violent crime rate and changes in greenhouse gas emissions




The results from each of our fact-checking manipulations are shown in Figure 2. Respondents who said the politician’s claims were definitely or probably true are labelled as “incorrect,” and those who thought they were definitely or probably false are labelled as “correct.” The fact-checks were once again successful in informing some respondents. For the crime rate question, 55% of respondents who did not receive the fact-check answered incorrectly, compared to 43% who did—a difference of 12 percentage points. For the question about emissions, the share of respondents answering incorrectly was almost 18 points lower among those who received the fact-check than among those who did not.

Finding 2: Our fact-checks were once again able to modestly correct our respondents’ beliefs about the violent crime rate and Canada’s greenhouse gas emissions.

Pooling our two corrections together allows us to compare how respondents react to fact-checks from the three different types of sources in this experiment. Respondents’ overall accuracy was scaled from 0-100, where 100 is the most accurate, as in respondents rated both claims as definitely false. The average estimated correctness of the respondents for each fact-checking source is displayed in Figure 3.

FIGURE 3. Average Effectiveness of a Fact-check by Source



As expected, the journalist was the most successful at correcting misperceptions. Respondents' average accuracy was 8 points higher when exposed to the journalist fact-check (41 on the 0-100 scale) than when they received no fact-check (33). However, the fact-checks from both the politician and the social media user also significantly improved our respondents' accuracy by 6 points each (39 vs. 33), such that there was no statistically significant difference between the effectiveness of fact-checks from any of the three sources. Respondents were not as responsive to the source of the fact-check as we might expect given our previous respondents' sharply differentiated evaluations of the credibility of different sources.

Finding 3: Averaged across our two fact-checking studies, all three of our fact-checks successfully corrected respondents. The journalist fact-check, however, was not any more effective than those from a Twitter commenter or a politician. Respondents were not particularly responsive to the source of the fact-check.

FINDINGS: TRUST AND INTERNAL EFFICACY

There has been a large volume of academic research on the effectiveness of fact-checks in the United States.⁴ This literature has generally shown them to be effective at correcting public misperceptions, if not entirely able to change higher-order beliefs about candidates or policy.⁵

However, there has been far less research on other downstream implications of fact-checking on public attitudes. First, fact-checking might undermine the trust people have in information provided by politicians or other representatives of public institutions, who are often the target of fact-checks.⁶ Second, fact-checking may influence citizens' *internal efficacy*—or their belief that they can understand and participate in politics. On the one hand, it is possible that fact-checks can improve efficacy helping members of the public feel more informed; on the other, they may be discouraging, particularly for individuals whose beliefs are being corrected.

In our most recent survey, we randomly exposed respondents to a series of four fact-checks of MPs by journalists. Once again, all actors in these fact-checking scenarios were fictional. This time, however, we asked respondents about their beliefs on the facts that were being corrected beforehand. These questions were about the violent crime rate (reality: lower now than 10 years ago), greenhouse gas emissions (reality: lower in 2017 compared to 2015), refugee levels (reality: lower now compared to 2016), and the unemployment rate (reality: lower now than in 2015).

Respondents could either get anywhere between zero to four fact-checks. After the fact-checks, we asked questions related to their trust in politicians as a source of information, and their belief in their ability to understand and participate in politics—or internal efficacy.⁷ Our scales for trust in politicians and internal efficacy were both constructed to run from 0-100, where 100 is the most trusting or efficacious.

⁴ Wintersieck, A. L. (2017). Debating the Truth: The Impact of Fact-Checking During Electoral Debates. *American Politics Research*, 45(2), 304–331. Wood, T., & Porter, E. (2019). The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence. *Political Behavior*, 41(01), 135–163.

⁵ Nyhan, B., Porter, E., Reifler, J., and Wood, T. (2019). Taking fact checks literally but not seriously? The effects of journalistic fact-checking on factual beliefs and candidate favorability. *Political Behavior*. First View. <https://doi.org/10.1007/s11109-019-09528-x>


⁶ Foa, R. S., & Mounk, Y. (2016). The danger of deconsolidation: The democratic disconnect. *Journal of Democracy*, 27(3), 5-17.

⁷ Respondents were asked to extent to which they trusted or distrusted politicians to provide objective, factual information (Completely trust to completely distrust, 5-point). They were also asked three questions tapping into internal efficacy: 1) I think that I am better informed about politics and government than most people; 2) I consider myself to be well qualified to participate in politics; and 3) I feel that I have a pretty good understanding of the important political issues facing the country (Strongly agree to strongly disagree, 5-point). The purpose of the experiment was masked such that these questions were mixed among others evaluating respondents' external efficacy towards political institutions and their trust in the media.

First, we find no evidence that fact-checks reduce the trust people have in information provided by politicians. Those who received no fact-checks had an estimated score of 44 out of 100, which dropped to 41 among those exposed to four fact-checks, but this difference is not statistically significant ($p \sim 0.15$).

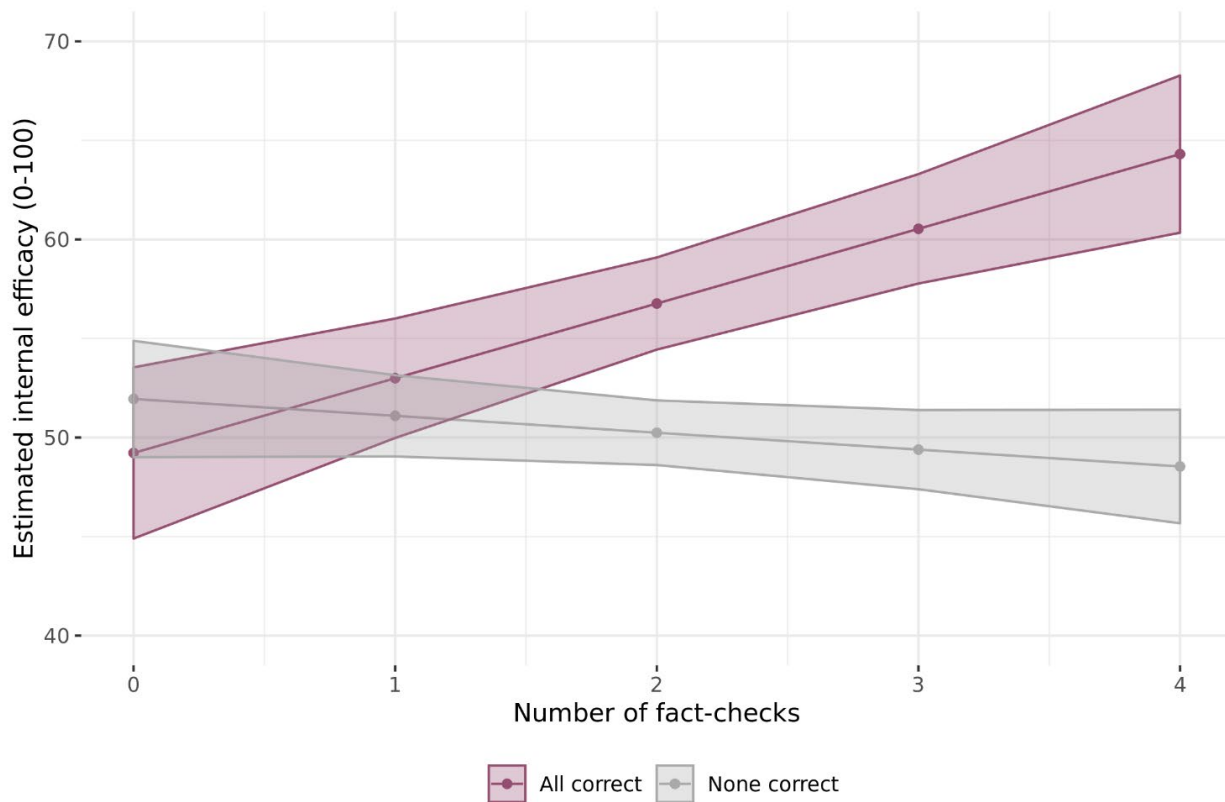
Finding 4: We find no evidence that exposure to fact-checks of politicians reduces trust in politicians as a source of political information.

FIGURE 4. Estimated efficacy across number of fact-checks for all respondents



There does, however, appear to be an important effect of fact-checking on internal efficacy. Fact-checking appears to *improve* internal efficacy on average. Respondents who were not exposed to fact-checks reported an internal efficacy score of 51 on the 0-100 scale. This improved to 55 among respondents exposed to all four fact-checks, which is a statistically significant difference ($p < 0.02$). This result is plotted in Figure 4 above.

FIGURE 5. Estimated efficacy across number of fact-checks for those who answered correctly or incorrectly across all questions



However, this average masks a significant amount of variation in the effect of fact-checks on efficacy. Specifically, it appears that the effect of fact-checking depends on how right or wrong a respondent is on the subjects being fact-checked. As seen in Figure 5, among those who responded correctly to all four questions at the outset, efficacy is 12 points higher for those who received all four fact-checks (64) compared to those who received none (52). By contrast, those who answered all four questions incorrectly before the fact-checks saw no significant change in their efficacy in either direction when they received more fact-checks (49 vs. 52). While fact-checking improves internal efficacy on average, the increase is particularly pronounced for those who have their beliefs re-affirmed rather than those whose beliefs are corrected by fact-checks. One consequence of this is that fact-checking increases the inequality in internal efficacy between those with high and low levels of policy knowledge.

Finding 5: Our evidence suggests that fact-checking improves individuals’ belief in their capacity to understand and participate in politics, particularly among those who have their correct beliefs reinforced by fact-checks. There is no evidence of a corresponding decrease in efficacy among those who had incorrect beliefs on the subject of the fact-check.

CASE STUDY

TRUDEAU AND BLACKFACE

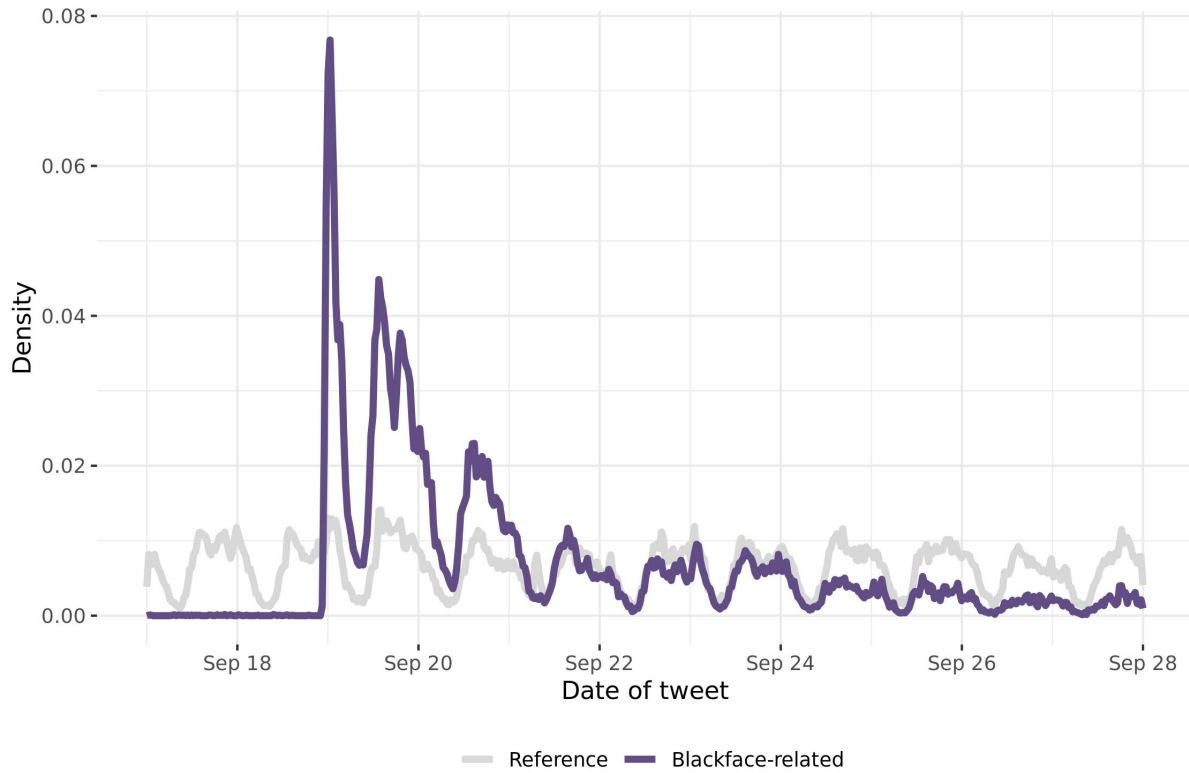
On Sept. 18, Time magazine published a photo of Prime Minister Justin Trudeau from a 2001 yearbook from the private school where Trudeau taught at the time. The school had staged an *Arabian Nights*-themed gala and Trudeau had dressed as Aladdin, wearing a turban with his face darkened. By the following morning, more photos or videos had surfaced showing Trudeau in brownface or blackface on two other occasions.

When a political story emerges as suddenly as this one did, it allows us to observe how social media behaviour changes in response. Here we seek to understand the reach of the story, the subsequent flurry of social media activity, and the extent to which the parties and their supporters engaged on the issue.

FINDINGS: EMERGENCE AND DECLINE OF BLACKFACE HASHTAGS

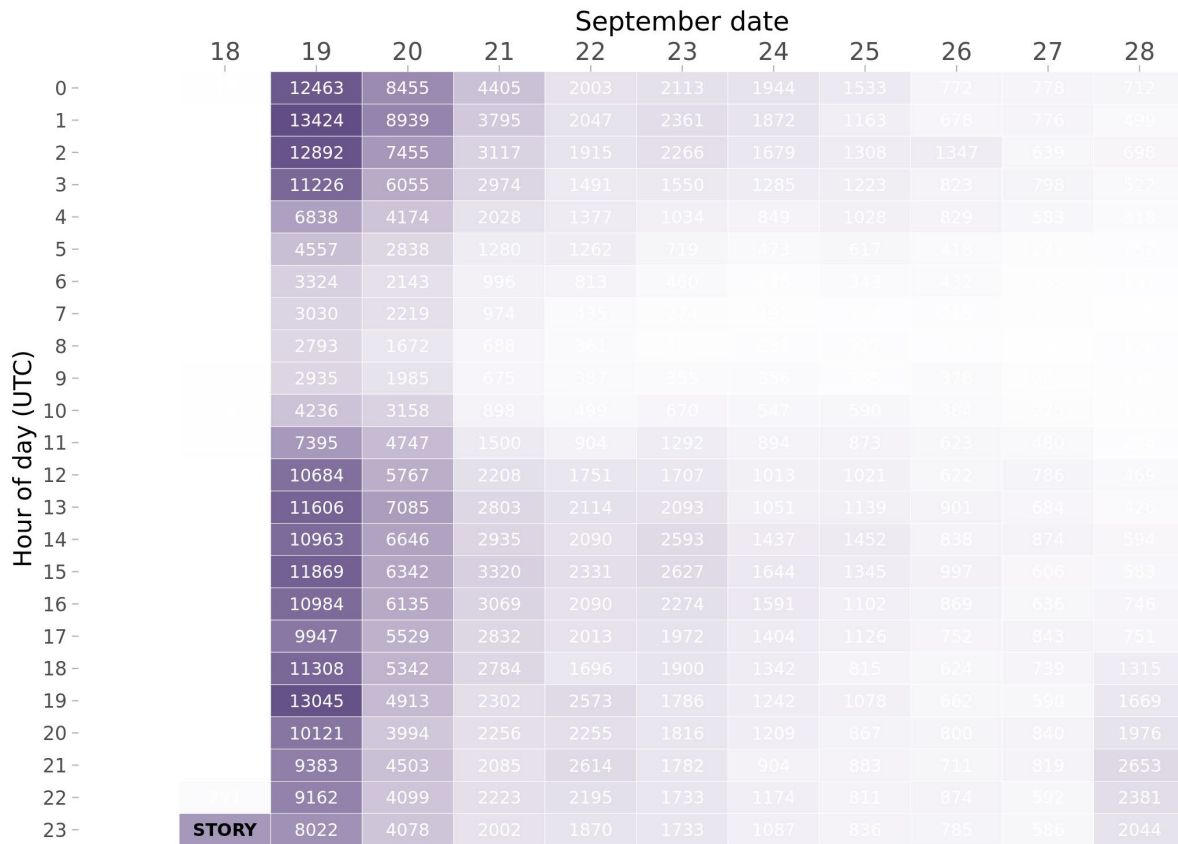
Figure 6 below shows the rapid emergence and decay of Twitter activity on the blackface controversy. Here we look at every tweet that contains hashtags, keywords or URLs associated with the story. The density plot shows the explosion of blackface-related language just after the story broke on the evening of Sept. 18 (all times in UTC), and the subsequent flurry of discussion afterwards. Indeed, the number of tweets was highest just after the story broke, with less activity overnight and a rebound the next day. For comparison, we included general election-related tweets, which have a relatively stable density throughout the campaign. This mapping strongly suggests that the online discourse around the controversy declined dramatically by Day 3.

FIGURE 6. Density plot of tweets including a hashtag, URL or keywords for blackface-related content as compared to general-election related content




Another view is a heatmap found in Figure 7. Again, we see the rapid emergence of the story and the steady decline over the subsequent week. While discussion continued to be very high for two to three days, the heatmap shows how quickly the Twitter population lost interest.

FIGURE 7. Heatmap of tweets including a hashtag, URL or keywords for blackface-related content, labelled with number of tweets for that day-hour




The rapid decay of such an explosive story is worth noting. By one week after the story broke, there was very little general Twitter interest in the story. The pattern on Facebook is similar, as shown in Figure 8, although the volume of public posts is much smaller.

FIGURE 8: Daily count of public Facebook posts including reference to blackface-related content



While the general population quickly lost interest in discussing the story on Twitter, both journalists and political candidates may have had incentives to continue engaging with the controversy on the platform. We collected all tweets from major-party candidates, and from approximately 700 journalists and 200 media outlets, that either (a) included the words “blackface” or “brownface” or related hashtags, or (b) linked to a story about the controversy. As with public Twitter users, news outlets and journalists demonstrated a similar spike in activity immediately after the story broke, with a large decay in subsequent days. There is not a noticeable difference in the decay between journalists and the general public—their interest in the story waned at a similar pace.

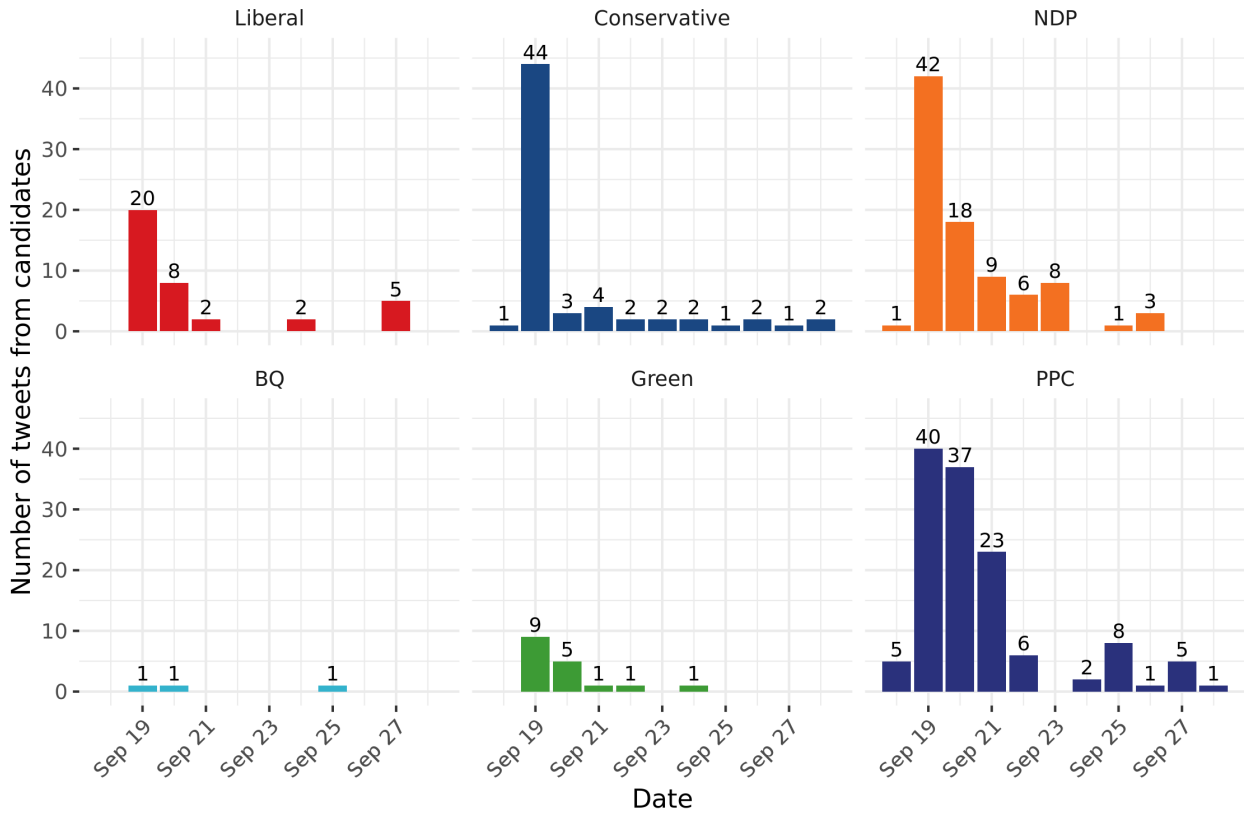
FIGURE 9. Daily count of Twitter posts from journalists that include reference to blackface-related content



Non-Liberal candidates may have had stronger incentives to keep the story alive. However, like journalists and the general public, their tweets on the controversy also dwindled a few days after the story broke. Figure 10 shows the number of original tweets by candidates on the blackface story. Here we see that NDP, People’s Party of Canada and Conservative candidates engaged the most on this issue, with Liberal, Green and Bloc Québécois candidates having minimal engagement. We also see another steep decay in interest after a few days of the parties heavily pushing the issue. Interestingly, Conservative candidates—on Twitter, at least—have not heavily focused on the issue except on the day after the story broke.

Finding 1: Social media discussion of the blackface controversy spiked the day after the story was reported, but declined dramatically by Day 3. Journalists and news outlets, candidates and the public all showed similar levels of decline in interest.


FIGURE 10. Number of tweets on blackface-related hashtags by candidate accounts



FINDINGS: PARTISAN DIFFERENCES AND ISOLATION

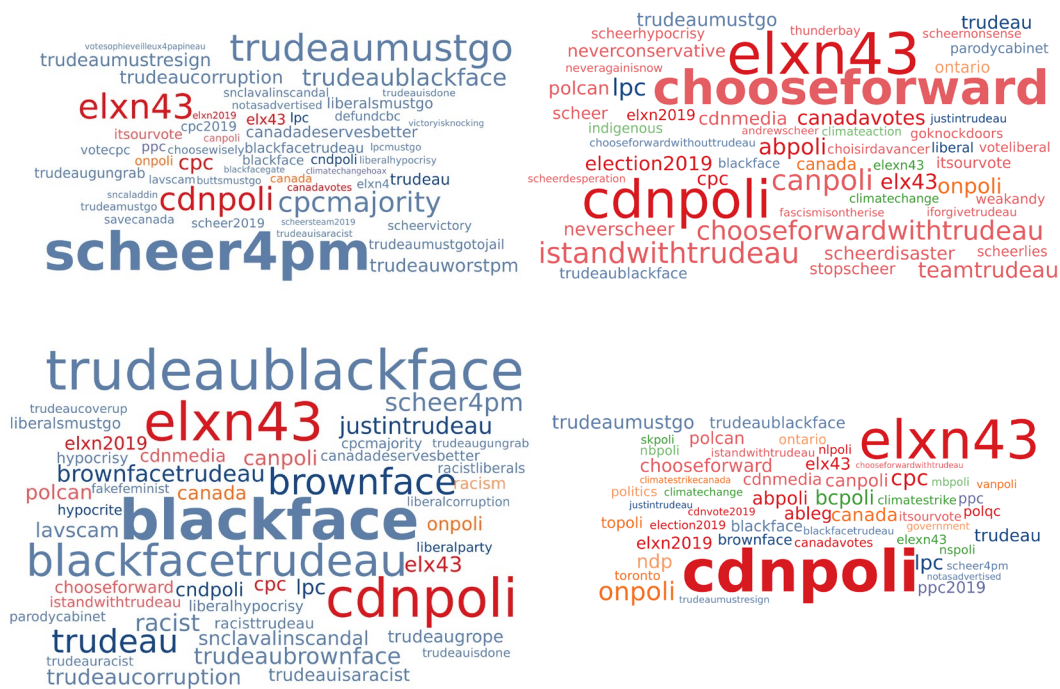
Next, we look at who among the broader Twitter population is engaging with the story and associated hashtags through tweeting or retweeting. Here we identify likely partisans by looking at the major-party candidates they follow. We identify approximately 55,000 likely partisans and examine their behaviour across relevant hashtags to get an overall understanding of how different partisan groups are engaging with the story. Figure 11 compares activity of partisan groups on hashtags related to the story (e.g. #brownface, #blackface) with general election hashtags (e.g. #elxn43, #cdnpoli), commonly used right-wing hashtags (e.g. #scheer4pm, #trudeaumustgo, #ppc2019), and commonly used left-wing hashtags (e.g. #chooseforward, #istandwithtrudeau, #ndp). Liberal, NDP, Bloc Québécois and Green supporters all have minimal engagement with the blackface-related hashtags. PPC supporters are not engaging heavily with the hashtags, either, instead continuing to focus on more general right-wing hashtags. It is, in fact, Conservative supporters who have been pushing this messaging, with their engagement on the blackface-related hashtags passing even that of their more commonly used pro-right hashtags.

FIGURE 11. Percentage of partisans tweeting or retweeting blackface-related hashtags as compared to general election hashtags and common left- and right-wing hashtags



Is this messaging reaching beyond Conservative supporters? One way that individuals break out of their communities on Twitter is through the use of hashtags—anyone can see what any other user has posted using a given hashtag. Here, we look at how often certain election-related hashtags coincide with other relevant hashtags. Figure 12 shows co-occurrence word clouds of select hashtags. The bolded hashtag is the anchor hashtag with which frequencies were compared. The colour is based on the partisan-congeniality of the hashtag, with a solid colour indicating more than 50% of the partisan users on that hashtag are drawn from that party, and a pale colour indicating that more than half of the users come from other parties. A more colourful word cloud indicates that a wider variety of partisans are engaging with the hashtag. For example, #cdnpoli is used more by Liberals than anyone else; however, supporters of other parties are also posting on this “contested” tag. Conversely, the #chooseforward and #scheer4pm hashtags reach primarily Liberal and Conservative partisans, respectively.

FIGURE 12. Hashtag co-occurrence word clouds with anchor hashtag bolded and all hashtags coloured based on the partisan group that uses the hashtag most



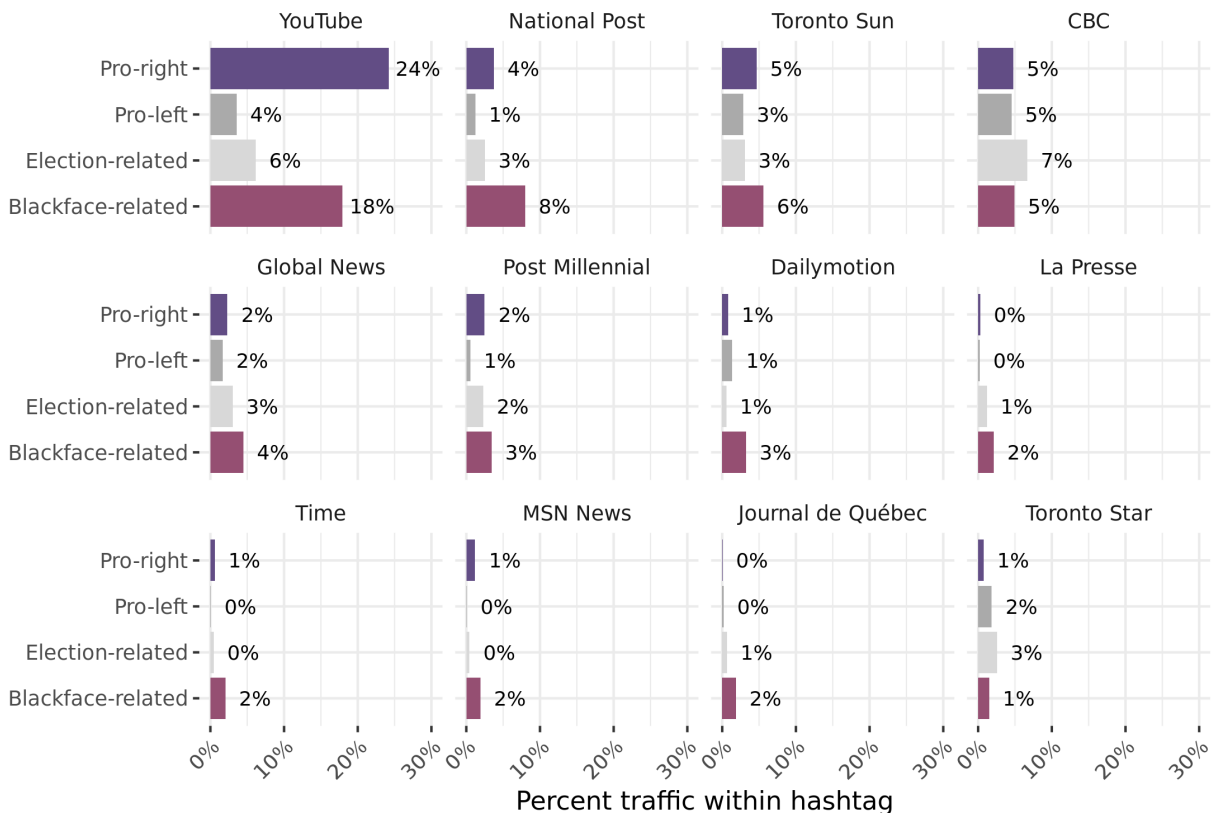
The #blackface hashtag is also on the homogenous side, with mostly Conservative and PPC partisans exposed to that messaging. Not only is the #blackface hashtag being primarily used by Conservatives, but it is being used in a way that does not reach partisans from other parties.

Finding 2: Among partisan Twitter users, Conservatives are driving the conversation about the controversy. The blackface-related hashtags are disproportionately populated by right-leaning partisans who are largely speaking among themselves.

FINDINGS: OPPORTUNISM AND MANIPULATION

New hashtags are often susceptible to opportunism, where individuals try to use a trend to promote unrelated content, but that does not appear to be a significant problem here. We examine the external sources being linked from four groups of hashtags: election-related, blackface-related, pro-left and pro-right. Figure 13 shows the 12 domains most frequently linked from these hashtags as a proportion of total number of links shared on the hashtag. There is a large number of YouTube videos being linked from both blackface-related and pro-right hashtags. Unsurprisingly, due to the original Time article that broke the story, links to Time are more common than otherwise expected in discussions of Canadian politics. Beyond these two differences, however, the rest of the top domains are mainly Canadian media sites, with only Dailymotion standing out as a clear exception. It does not appear that unusual content is being shared or that this hashtag is being used to drive users to more obscure websites.

FIGURE 13: Links to external websites from select Canadian politics hashtag groups. Top 12 websites shared on blackface-related hashtags are shown

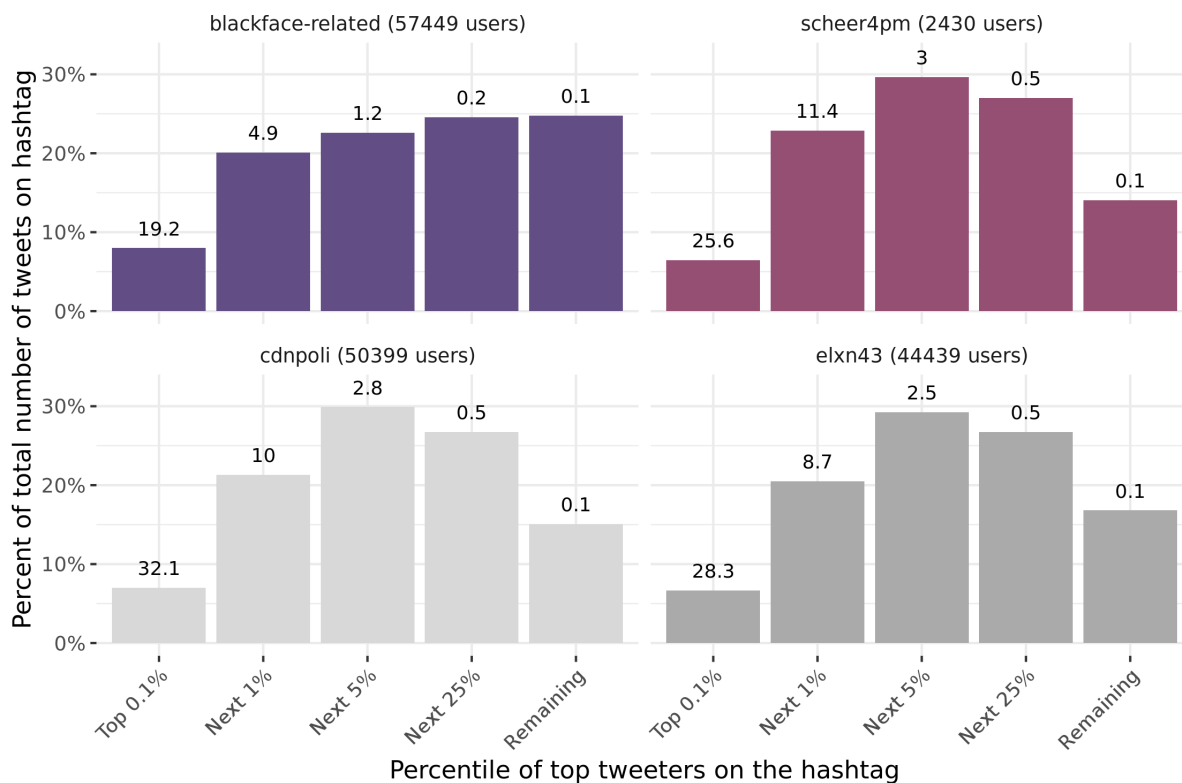


A second consideration is that, due to their trending status and high visibility, new trending hashtags attract spammers and/or bots that amplify the hashtag and attempt to push a particular narrative. Here, we briefly

look at the extent to which activity on this hashtag is driven by automated or human accounts tweeting at a high frequency. To test this, we compare post frequencies on the blackface-related hashtags against general hashtags with a longer lifespan to see if the activity is driven by a relatively small number of users. While the top 0.1% of accounts accounted for about 7% of blackface-related hashtag tweets (a considerable 230 tweets per user over a 10-day period), the distribution of tweet frequency is similar to comparable hashtags. By this metric, at least, it does not appear that the immediate popularity of these hashtags was driven by what is increasingly called “coordinated inauthentic behaviour.”

Finding 3: The blackface-related hashtags do not appear to be disproportionately populated by inauthentic activity.

FIGURE 14. Concentrated activity across hashtags from Sept. 17-28, with numbers above bars indicating average number of tweets per user per day in each category



APPENDIX

Our survey data team conducted an online panel survey of 1,594 Canadian citizens 18 years and older using the sample provider Qualtrics. The sample was gathered from Sept. 19-24. Data was weighted within each region of Canada by gender and age based on data from the 2016 Canadian census. We used an iterative proportional fitting algorithm for our weighting procedure with a minimum weight of 0.34 (N=24) and a maximum weight of 1.52 (N=9).

Survey respondents were asked questions related to basic demographics, as well as their partisan, ideological and issue preferences. They were also asked to identify their recent exposure to the news media. The median time it took respondents to complete the survey was 24 minutes. The survey instrument is available upon request. We present 90% confidence intervals for each of our figures. Our analyses on the determinants of support for fact-checking can be generalized to the population +/- 0.01-0.02 with 90% confidence.

Facebook data is sourced from CrowdTangle, a social media analytics tool owned by Facebook. CrowdTangle tracks public posts on Facebook, Instagram and Reddit, made by public accounts or groups. The tool does not track every public account and does not track private profiles or groups, so this data is not representative of performance across the entire platform. The numbers shown here reflect public interactions (likes, reactions, comments, shares, upvotes and three-second views), but do not include reach or referral traffic. It does not include paid ads unless those ads began as organic, non-paid posts that were subsequently "boosted" using Facebook's advertising tools. Because the system doesn't distinguish this type of paid content, note that some high-performing content may have had paid distribution. CrowdTangle also does not track posts made visible only to specific groups of followers.

**DIGITAL
DEMOCRACY
PROJECT**

