

Social Distancing as a Public Goods Dilemma: High Economic Cost Reduces Voluntary Compliance¹

Eric Merkley
University of Toronto
eric.merkley@utoronto.ca

Peter John Loewen
University of Toronto
peter.loewen@utoronto.ca

Abstract

Participation in social distancing can be seen as a contribution to a public good that is influenced by 1) the marginal costs and benefits of those contributions and 2) expectations that other citizens will participate. We test our theory using an official government economic report on job loss as an exogenous, negative information shock. Using Canadian data, we show that this shock increased aggregate-level mobility and reduced self-reported social distancing among respondents surveyed throughout the pandemic (N=17,539), especially for younger respondents – a group that faces higher costs relative to benefits of compliance. We also conduct three survey experiments on nationally representative samples to unpack a possible mediating effect of expectations of others' participation. Our results reinforce our principal findings, while also showing that 1) information on prospective economic cost reduces expectations of compliance by other citizens; and 2) expectations of compliance by others cause expectations of respondents' own compliance.

¹ We grateful for the helpful feedback from James Fowler and Dominik Stecula. Thanks as well to the rest of the Media Ecosystem Observatory research team: Taylor Owen, Derek Ruths, Aengus Bridgman, Oleg Zhilin, Sean Nossek, Lisa Teichmann, Elisa Chaudet, and Julia Ma. This work is funded by the Department of Canadian Heritage's Digital Citizens Initiative.

The COVID-19 pandemic has posed a challenge to governments and citizens unlike any in recent history. Containing its spread requires citizens to adhere to the advice of epidemiologists and public health experts to engage in social or physical distancing. Among other actions, this behavior involves maintaining space in public, reducing unnecessary trips outside the home, avoiding workplaces, and often limiting one's in person social interactions. Practicing social distancing as much possible protects individuals and prevents transmission within the community, but at a cost of serious social and economic disruption. Government policy – such as the closure of non-essential businesses and public spaces – is essential in maintaining compliance with social distancing. But successful containment also hinges on individual citizens adhering to these guidelines voluntarily. Maintaining compliance for the extended period of time required to develop and mass produce an effective vaccine is thus an enormous public policy challenge.

The social and natural sciences have spent much of the last century understanding how social dilemmas – where individuals' pay a cost to produce a good that society enjoys – can be resolved. The difficulty of studying social dilemmas in the real world is confounded by the presence of institutions (North, 1991), multiple forms of punishment (Henrich et al., 2006), and norms tailored to certain dilemmas over a long period of time (Ostrom, 2000). For example, consider the case of taxation. Most individuals in developed states comply by paying their owed taxes. This may be a result of a strong enforcement mechanism by the state, of citizens seeing that states credibly deliver on the goods promised in exchange for taxes, or because they have developed a series of norms around compliance with taxation independent of the likelihood of sanction. Likely, all of these factors matter, but estimating the effects of individual mechanisms in a fashion similar to that in laboratory experiments is very difficult, given the complexity of how social dilemmas are normally resolved in the real world.

We argue that public health during a pandemic is a public good to which individuals can make a costly contribution by engaging in social distancing. Because of the rapidity of the spread of the COVID-19 pandemic and the massive, sudden response of governments in imposing widespread restrictions and lockdowns, case-specific norms and accompanying institutions have not had a chance to evolve in support of resolving the social dilemma. For example, practices for when and how violations of social distancing are sanctioned by the police are still developing. Norms for how people respond to fellow citizens defecting from social distancing are not yet universal. As importantly, behavioral and economic strategies to cope with the costs of social distancing are not well-developed. The response of citizens to the pandemic thus provides a unique opportunity to examine a large-scale, N-person public goods game.

In the context of COVID-19, the decision to contribute is thus principally a function of its costs and benefits and other more basic responses to public goods dilemmas (Rand, Greene, & Nowak, 2012). From this logic, we argue that high economic costs are likely to reduce compliance with social distancing and expectations of compliance by other citizens. Furthermore, we expect that effects on compliance will be stronger among citizens who face a disproportionately higher share of those expected costs relative to benefits – principally the young. We further anticipate partial mediation, where the effect of economic cost on compliance is mediated by changes in their expectations of other people's behavior.

Our paper provides two sets of analyses to test our theory. First, we use the Canadian April jobs report as an exogenous, negative economic information shock that is expected to reduce social distancing compliance. We leverage individual-level, cross-sectional survey data collected before and after the release of the report, along with panel-based survey responses with re-contacts occurring before and after the release of the jobs report, and aggregate-level mobility data to test whether the information shock reduced compliance. We find consistent evidence that it did reduce compliance with social distancing, especially among younger respondents.

Second, we present the results of three experiments conducted on nationally representative samples. We show that exposure to information about high prospective economic costs of social distancing is associated with: 1) lower expectations of social distancing compliance by other citizens; and 2) lower expectations of one's own compliance, especially among younger respondents. We also show 3) that expectations of social distancing compliance by other citizens affect expectations of one's own compliance. The implication is that maintaining compliance with social distancing amidst economic collapse may become much more difficult as economic costs become greater.

Social Distancing and Public Goods Co-production

Containment of the COVID-19 pandemic requires citizens to engage in widespread social or physical distancing on the advice of epidemiologists and other health experts. Understanding what groups of citizens are likely to take heed of expert advice is thus vitally important. Many scholars see the central challenge as dealing with individual-level predispositions often at odds with expert messages, such as partisanship (Merkley & Stecula, 2020; Pasek, 2017; Tesler, 2018), ideology (Campbell & Kay, 2011; Lewandowsky & Oberauer, 2016), cultural values (Kahan et al., 2011; Leiserowitz, 2006), and anti-intellectualism (Merkley, 2020; Motta, 2018) that cannot be overcome by simply providing citizens with more information from experts. We have seen this problem in the United States where attitudes about, and behaviors related to, COVID-19 are heavily structured by partisanship (Allcott et al., 2020; Cornelson & Miloucheva, 2020; Gadarian et al., 2020).

This approach is valuable, but its emphasis on understanding *attitudes* have some important limitations when applied to the understanding the *behavior* of citizens. Citizens tend to form attitudes in line with their social identities because motivations toward direction can dominate accuracy in politics (Kunda, 1990; Lodge & Taber, 2013). They do not face costs for holding inaccurate beliefs, while they gain social-psychological benefits for supporting their social groups. Individual decisions to engage in social distancing, however, have clear costs and benefits for themselves that also partially spill over to other citizens. They are behaviors with consequences.

Public health provision, such as the sanitation, clean air and water, and widespread vaccination, has some of the characteristics of a public good. It is non-excludable and mostly non-rivalrous – as in the number of consumers of the good do not diminish the good (Fisman & Laupland, 2009). However, provision of this good can be partially characterized by its co-production by governments and citizens (Ostrom, 1996). Public health depends on the readiness of citizens to contribute, often at the instruction of public health officials. For example, herd immunity is partially dependent on the willingness of citizens to vaccinate (Siegal et al., 2009). Public health during a pandemic is similarly dependent on the positive externalities generated by citizens voluntarily engaging in social distancing. Government policy, of course, plays an essential role by closing non-essential businesses

and public places to limit transmission, but in a democracy some level of public health production – likely a high level – is going to be dependent on the voluntary actions of individual citizens and their willingness to comply with health directives.

Consequently, the level of public health co-provision is going to be partially determined by the costs and benefits of participation for citizens, especially in the absence of strong norms towards compliance or cooperation. For example, some scholars have argued that disease elimination through voluntary childhood vaccination is likely impossible because parents will rationally refuse to vaccinate when the risk of infection is less than the risk of vaccine side-effects – they free-ride off of herd immunity (Bauch & Earn, 2004). The maintenance of public health during a pandemic also depends on the costs and benefits of contributions for individual citizens. Marginal increases in costs relative to benefits are likely to decrease participation.

This is a problem because social distancing entails enormous short and long-term economic costs on society (Fenichel, 2013). As citizens face the high economic costs of social distancing, we should expect reduced compliance. This pattern is not unique to pandemics. For instance, scholars have shown that the willingness of citizens to adopt pro-environmental attitudes and behaviors is shaped by economic conditions and the costs and benefits of policies (Aldy et al., 2012, Aklin & Urpelainen, 2013, Ansolabehere & Konisky, 2014; Brulle et al., 2012; Scruggs & Benegal, 2012).

During the COVID-19 pandemic, the crushing economic costs of voluntary social distancing and government lockdowns were strongly brought to the forefront of political discussion with monthly releases of unemployment data. We treat these releases as exogenous information shocks that inform citizens of the economic costs of social distancing. We focus on the April jobs report in Canada, released on May 8, 2020, that revealed the the unemployment rate had increased to over 13% in April with 2 million jobs lost – by far the worst single monthly jobs report in Canadian history. We marshal two types of data to illustrate the effect of this exogenous shock on social distancing practice: individual-level survey data in which respondents self-reported their social distancing behavior and Google mobility data through April and May. We expect the April jobs report produced a reduction in social distancing practice.

H1: The April jobs report information shock is associated with less social distancing compliance.

Modeling compliance with social distancing needs to take into account that some groups of citizens receive comparatively less benefit from social distancing than others. Perhaps most notably, older citizens are far more at risk of serious complications from COVID-19 than those who are younger (Verity et al., 2020). The economic fallout of the pandemic is also especially devastating for younger citizens, dramatically reducing income and opportunity in prime earning years. This was also true of the Financial Crisis and Great Recession of 2008, but it is made especially acute in this case since younger citizens make up a disproportionate share of the devastated service sector (Lowrey, 2020). We expect the April jobs report to have had stronger effects on younger respondents.

H2: The effect of the April jobs report information shock on social distancing compliance is stronger among younger respondents.

We further unpack the mechanisms undergirding the relationship between economic cost and social distancing compliance. We expect not only a direct effect of economic cost on social distancing compliance, but an indirect effect through changes it generates in expectations of other people's behavior. If people believe other citizens will fail to comply with social distancing guidelines, their own contribution to the public health is less likely to be consequential, and so they will fail to contribute. Scholars have used this logic to explain why, for instance, public support for climate change mitigation policy is higher when more countries participate (Bechtel et al., 2017; Bechtel & Scheve, 2013). This is also supported by theories of norms that would expect behavioral change in response to perceived social norm erosion (Akerlof & Kranton, 2000; Bicchieri, 2006; Pickup et al., 2020; Sinclair, 2012; Suhay, 2015).

To this end, we conducted a pair of experiments to examine the effect of exposure to economic information on expectations of compliance by other citizens (Experiment 1) and by oneself (Experiment 2). Experiment 1 randomly assigns projections of economic consequences of social distancing that vary in their severity and evaluates the effect of this intervention on expectations of social distancing compliance by other citizens. Experiment 2 randomly assigns projections of high economic cost and examines its effect on these expectations, as well as expectations of one's own behavior. These experiments took place before the economic consequences of the pandemic came into full view. We test the following hypotheses, where H4 and H5 directly map onto our hypotheses related to the April jobs report:

H3: Exposure to information about high prospective economic cost is negatively associated with expectations of social distancing compliance by other citizens

H4: Exposure to information about high prospective economic cost is negatively associated with self-expectations of social distancing compliance.

H5: The effect of prospective economic cost treatment on self-expectations of social distancing compliance is stronger among younger respondents.

We were not able to exogenously manipulate expectations of the social distancing compliance by other citizens to produce causal evidence of mediation in the second experiment. So we conducted a third experiment where we randomly assign respondents polling information on the willingness of people to engage in a variety of social distancing practices. We do this to provide causal evidence of a link between expectations of other citizens' social distancing compliance and expectations of one's own behavior:

H6: Exposure to polling information about long-term social distancing compliance is positively associated with self-expectations of social distancing compliance.

The April Jobs Report as an Exogenous Information Shock

Economists and public policy experts expected the COVID-19 pandemic to have negative consequences for the national economy. Government lockdowns shuttered most non-essential businesses, while individual citizens limited their social and economic activity in order to "flatten the curve." However, typical indicators of economic performance are ill-equipped to shed light on such a rapidly unfolding situation. Canadians were in the dark on the full scope of the economic damage

until Statistics Canada released the results of its April labour force survey on May 8, 2020. This was the first complete picture of the economic consequences of the COVID-19 pandemic. The report revealed that 2 million jobs were lost over the course of April – the highest monthly figure on record – with young people bearing the brunt of the shock. The unemployment rate in Canada rose to 13%, which was likely an underestimate since recently employed people dropped out of the labour force as a result of the pandemic and are not included in the calculation of that statistic. Statistics Canada estimated the unemployment rate was closer to 18% after adjusting for this fact. This was the worst jobs report in Canadian history.

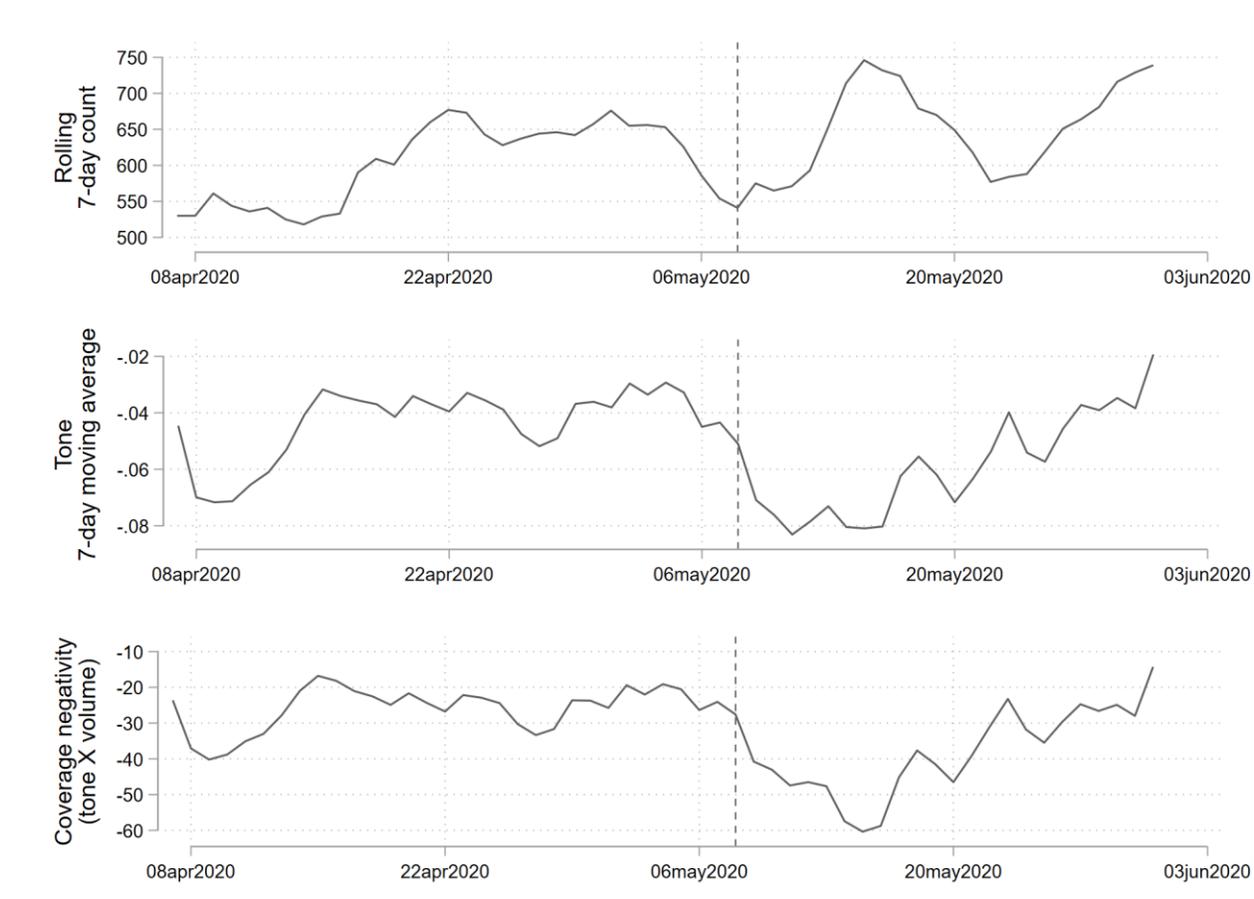


Figure 1. 7-day rolling sum of economic news content (top); 7-day moving average of headline sentiment (middle); composite coverage negativity index (bottom). Dashed line represents release of the April jobs report.

The report also precipitated a torrent of negative news about the economy. We downloaded all newspaper headlines between April 1 and May 27 from Canada’s two national circulating newspapers from *Lexis Uni* – the *Globe and Mail*, and the *National Post*, along with two major regional papers, the *Toronto Star*, and the *Montreal Gazette* (N=5,422). Headlines were selected if they contained the following keywords in the body of the article: (econom* or jobs or employment) and Canada. These words were used to restrict our search to economic news focused on Canada. The

top panel of Figure 1 contains a 7-day rolling sum of economic news articles from these four sources. The dashed line indicates the day the jobs report was released. A sharp spike in economic news coverage clearly coincides with the release of the report.

We can also measure the sentiment of our headlines. We use the economic sentiment dictionary of Hopkins et al. (2017) that counts negative and positive economic terms to evaluate tone in economic news content.² We apply this dictionary to our headlines and calculate a tone score per article by taking the difference between the positive and negative words, which we average daily. The middle panel of Figure 1 displays a 7-day moving average of tone in our economic news headlines, where the dashed line represents the release of the jobs report. As is clear, the release of the report coincides with a sharp and sustained drop in headline sentiment.

Finally, we construct a composite score of coverage negativity by adjusting our tone measure by changes in the volume of content (volume * tone). The intuition here is that negative sentiment will matter less when coverage is low. This measure is displayed in the bottom panel of Figure 1. Again, the jobs report release coincides with much more negative economic news. We take this as compelling evidence that the release of the economic jobs report provided an exogenous information shock.

Individual-Level Data and Design

The April jobs report provided an information shock to mass publics about the economic costs of controlling the COVID-19 pandemic through social distancing and government imposed lockdowns. We expect this shock reduced social distancing compliance (H1), especially among young people (H2). We test this by making use of surveys conducted by the Media Ecosystem Observatory (MEO). The MEO conducted seven surveys of approximately 2,500 Canadian citizens each over the course of April and May 2020. Quotas were set on gender, age, language, and region, and we further weight our data within region by age and gender using iterative proportional weighting to ensure sample representativeness. The fielding dates of each wave can be found in the supplementary information. We pool these respondents together (N=17,539).

In each wave of the survey we ask respondents whether or not they have taken a variety of actions over the previous week in response to the pandemic. We use four of these items that most closely reflect the social distancing guidelines provided by health experts (i.e. avoided bars, restaurants, and crowds; avoided in-person contact; kept a distance of two meters; avoided domestic travel) to provide a reliable index of this behavior (Cronbach's alpha = 0.75).³

² The keywords are as followings: fall*, recess*, problem, unemploy*, slump*, slow*, drop*, fear, debt, bad, plung*, jobless, loss*, bear, layoff* (negative); profit*, bull, growth, growing, grow, inflat*, invest* (positive). We replicate our main findings with the Lexicoder Sentiment Dictionary (Young & Soroka, 2012) as well, which can be found in Figure S2 in the online supplement.

³ The full battery asked respondents whether or not they took the following additional actions: 1) worked from home; 2) avoided grocery stores at peak times; 3) stocked up on provisions; 4) switched to virtual meetings; 5) switched to online shopping; 6) avoided public transit. We performed a PCA to extract two dimensions of social distancing behavior. One dimension is closely

We test H1 by estimating a model with robust standard errors that predicts our social distancing index with a daily trend and a dummy variable indicating whether the respondent was surveyed after the release of the jobs report ‘1’ or not ‘0’.

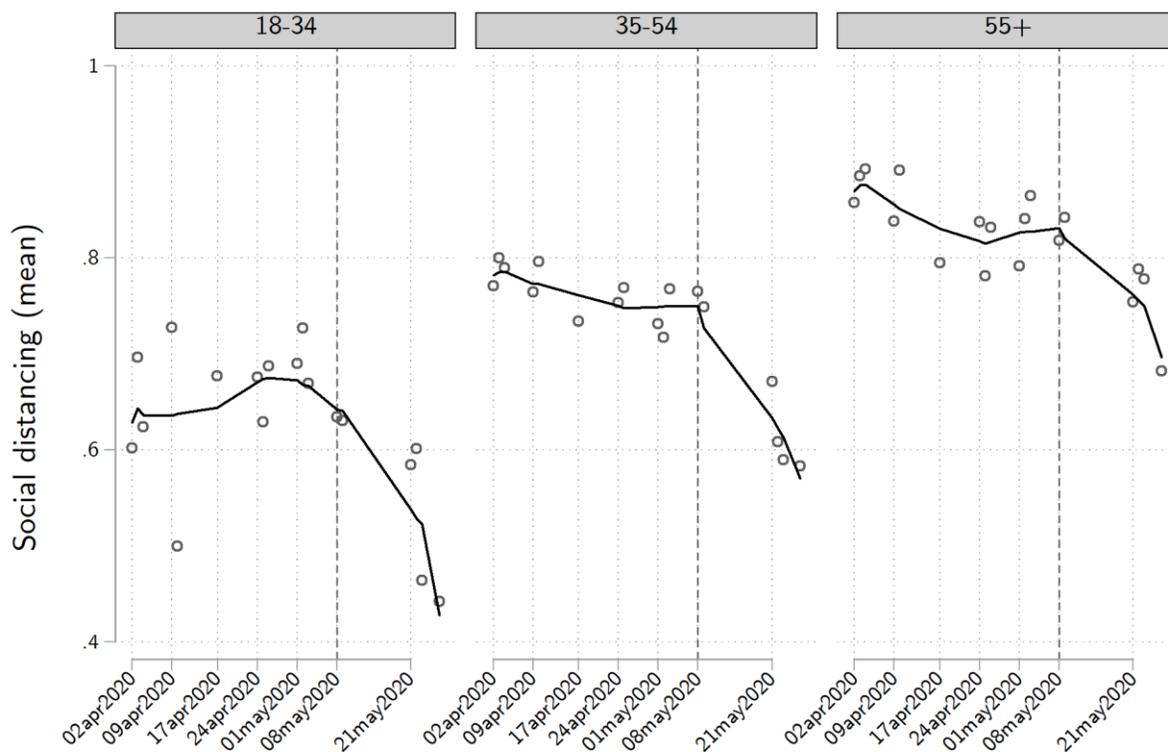


Figure 2. Mean social distancing by age group. Note: days with less than 100 respondents dropped from series. Lowess smoothed at a bandwidth of 0.5.

We evaluate the robustness of our finding by estimating models: 1) using an alternative index of social distancing extracted from a principal component analysis (labelled *Alternative*); 2) controlling for predispositions, news exposure, province-level re-opening policy, and demographics (*Controls*); 3) dropping respondents from the final wave conducted two weeks after the jobs report (i.e. wave seven; *Narrow 1*); and 4) using only the wave that overlaps the jobs report and the wave immediately prior (i.e. waves five and six; *Narrow 2*). The purpose of these latter two tests are to illustrate whether we still observe effects very close to the release of the report. We test H2 by interacting our dummy indicator with the age of the respondent. In order to increase confidence we are observing heterogeneous effects across age due to our theory, we control for interactions of the report release with COVID-19 risk perceptions, anti-intellectualism, ideology, news exposure, education,

related to our items of interests. We use this alternative measure of compliance for a robustness check.

urban/rural location, along with controls for gender and region. Descriptions of our variables can be found in Table S2 of the supplementary materials.

Figure 2 plots the average social distancing score through April and May. There appears to be a discontinuity in the series around the time of the jobs report, and this is most pronounced among those between the ages of 18 and 34.

We can also exploit the fact that around half of respondents in the first three surveys were re-contacted. Approximately half of waves four through seven were re-contacts of waves 1 through 3. As a result, some of these re-contacts occurred before the jobs report and some occurred after its release during the fielding of wave 5 ($N \sim 4910$). We test H1 by estimating the same models as above but with *changes* in social distancing as the principal outcome. We also include dummy variables for contact and re-contact waves. We test H2 by interacting our dummy indicator with the age of the respondent along with the same controls as the cross-sectional analysis as well contact and re-contact wave fixed effects.

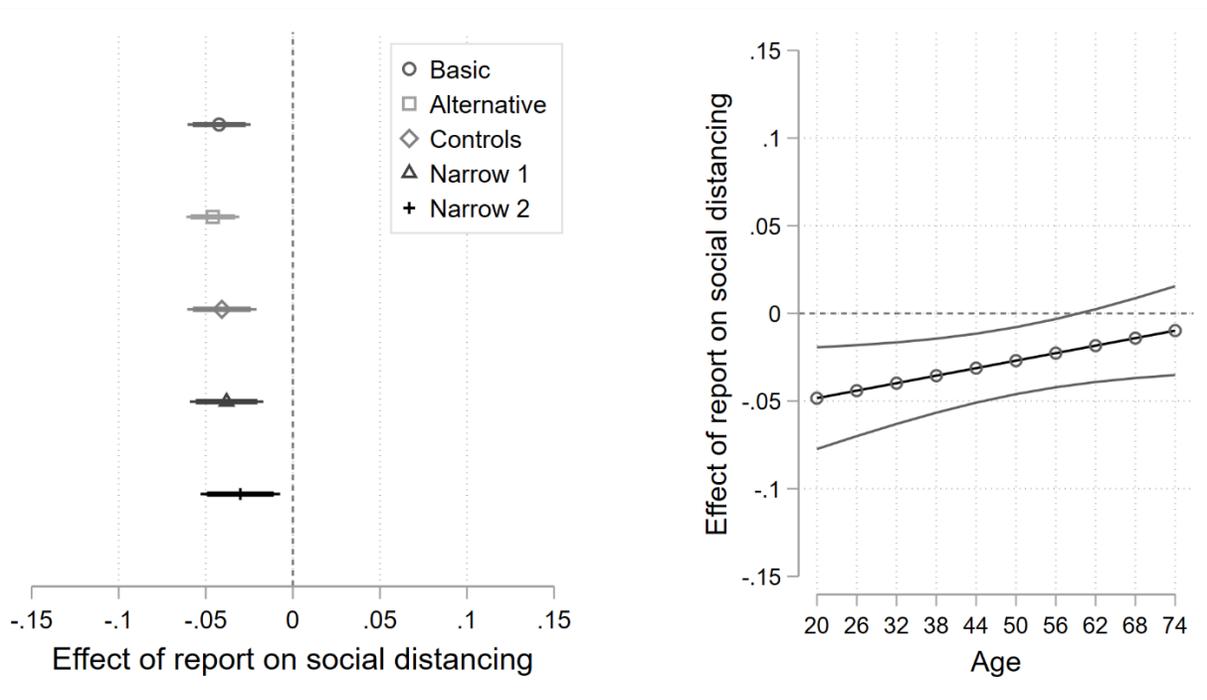


Figure 3. Effects of the job report release on social distancing compliance (left panel) across age group (right panel). Note: 95 and 90% confidence intervals in coefficient plot, 95% confidence intervals in marginal effects plot. ‘Alternative’ model uses social distancing outcome extracted from a PCA. ‘Controls’ model includes controls for province-level re-opening, COVID-19 risk perceptions, anti-intellectualism, ideology, news exposure, education, urban/rural location, gender, and region. ‘Narrow 1’ model drops respondents from the final wave, fielded from May 21-27, while ‘Narrow 2’ preserves only respondents in wave 4, immediately before the report release from May 1-5, and those in wave 5, surveyed from May 8-12.

Individual-level Results

Our cross-sectional results are presented in Figure 3. We find that being surveyed in the period after the release of the jobs report is associated with a 0.04 point reduction in self-reported social distancing compliance ($p < 0.001$) after controlling for a linear trend. This amounts to a 0.13 standard deviation or 5% decrease on this measure. This finding is robust to the use of an alternative outcome, the inclusion of controls, and narrower bands of analysis, including – remarkably – a comparison of respondents surveyed the day of the report’s release to the days after.

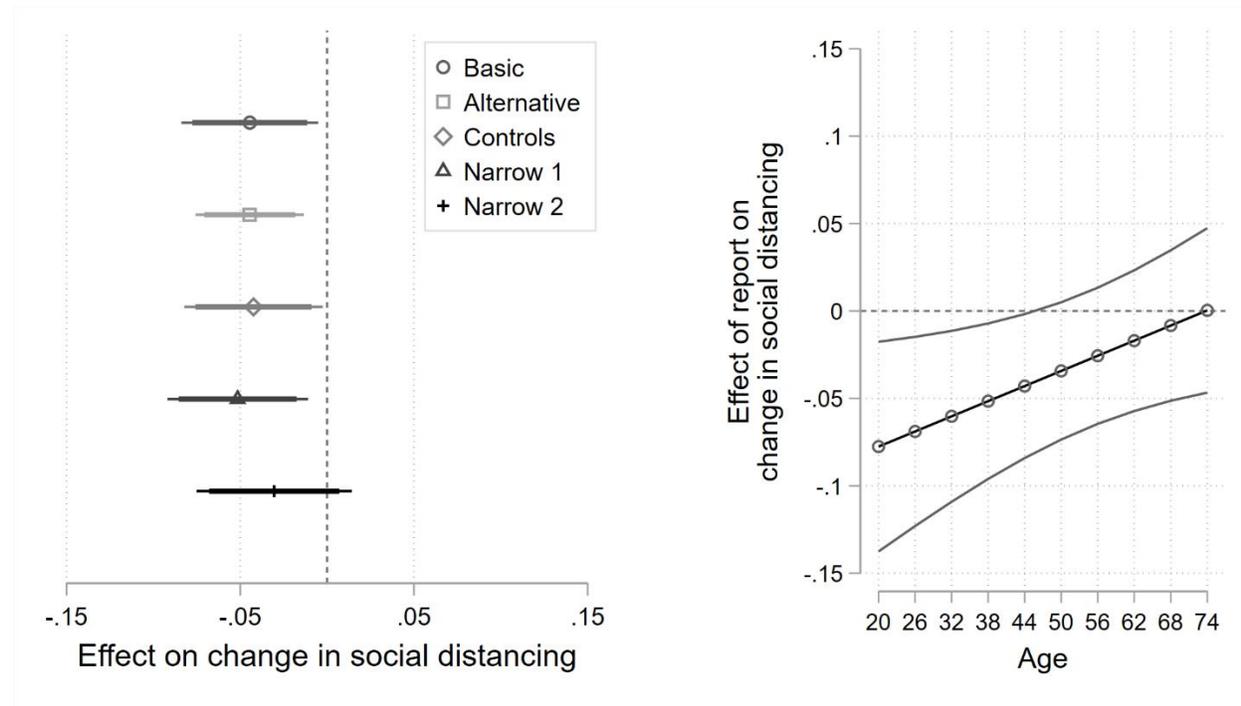


Figure 4. Effects of the job report release on changes in social distancing compliance (left panel) across age group (right panel). Note: 95 and 90% confidence intervals in coefficient plot, 95% confidence intervals in marginal effects plot. ‘Alternative’ model uses social distancing outcome extracted from a PCA. ‘Controls’ model includes controls for province-level re-opening, COVID-19 risk perceptions, anti-intellectualism, ideology, news exposure, education, urban/rural location, gender, and region. ‘Narrow 1’ model drops respondents re-contacted in the final wave, fielded from May 21-27, while ‘Narrow 2’ preserves only respondents who were re-contacted in wave 5, fielded from May 1-5 immediately before the report release, and those who were re-contacted in wave 6, surveyed from May 8-12.

This effect appears to be heterogeneous across age with the young most affected. The interaction between the treatment and age is significant ($p < 0.05$) and the effects are plotted in the right panel of Figure 3. Twenty-year-old respondents are expected to reduce their social distancing by 0.05 points in response to the jobs report or 0.16 standard deviations. This effect weakens as age increases and disappears for those over the age of 65. It is worth reiterating that this interaction effect is after controlling for the how social distancing varies before and after the report by respondents COVID-

19 risk perceptions, anti-intellectualism, ideology, news exposure, education, urban/rural location. Further, the interaction disappears when the model includes an interaction of the report release with a variable indicating respondents' level of unemployment risk of due to COVID-19, which in turn becomes significant, providing more evidence that the heterogeneous effects we observe are due to our theory rather than other factors that are also associated with age.

We find remarkably similar results with our panel data, displayed in the left panel of Figure 4. Recall that our panel respondents were initially contacted between April 2 and 19. Respondents who were re-contacted after the release of the jobs report (i.e. May 9-27) reduced their self-reported social distancing compliance by 0.04 points more than those who were re-contacted before the release of the report (i.e. April 24-May 8) after controlling for trending and wave fixed effects ($p \sim 0.03$). This finding is largely robust, though we lose statistical significance when analyzing only respondents re-contacted during waves 5 and 6 (*Narrow 2*). This is largely owing to a loss of power compared to the cross-sectional equivalent – the point estimate is similar.

Again, we find evidence of heterogeneous effects across age. Twenty-year-olds re-contacted after the release of the jobs report are expected to reduce their reported social distancing 0.08 points more than those re-contacted beforehand. This effect weakens until it becomes statistically insignificant for those 50 years and older. The interaction term is statistically significant ($p \sim 0.03$).

In short, we find evidence that self-reported social distancing significantly declined after the release of the jobs report. This occurred most sharply among younger respondents, and for reasons unrelated to their predispositions or other socio-demographic characteristics. We find this when comparing different survey respondents over time, and within-respondent changes observed before and after the jobs report release – compelling evidence in support of H1 and H2.

Google Mobility Data and Design

The above analysis has two key limitations. First, it hinges on the reliability of self-reported survey measures of social distancing compliance. Social distancing is a new concept in social science and much remains uncertain about to optimally measure this in self-reported survey measures. Recent work has shown that estimates may suffer upward bias due to social desirability (Daoust et al., 2020), though the gravity of this problem has been contested (Jensen, 2020), while others find that these measures nonetheless match behavior reasonably well (Gollwitzer et al., 2020). Second, our survey design was not one of a rolling cross-section, so we do not have continuous daily samples that would allow us to rule out time varying confounders. It is possible there are other changes happening around the time of the jobs report that produce less social distancing compliance – especially among young people. We mitigate this concern by showing effects that are very close to the release of the jobs report, but we cannot rule this problem out entirely.

Consequently, we supplement our findings with a time series analysis of Google mobility data from April 1-May 27 in Canada. We cannot examine heterogeneous effects using these data, and the level of granularity in the data does not allow us to do fine-grained spatial analysis. We can, however, evaluate whether there is a discontinuity in the mobility series corresponding to the release of the April jobs report after controlling for time varying confounders.

We create an index of three Google mobility series that we expect are more likely to be a function of voluntary compliance with social distancing guidelines: parks, retail and recreation, and groceries and pharmacy. We also construct an index of three series that are less likely to reflect voluntary compliance with social distancing, but rather occupation and employment patterns: residential (reverse-scaled), workplace; and transit mobility. A principal components analysis shows that these measures load together on two distinct factors as expected.⁴ Both of these measures are trend stationary according to Dickey-Fuller tests.

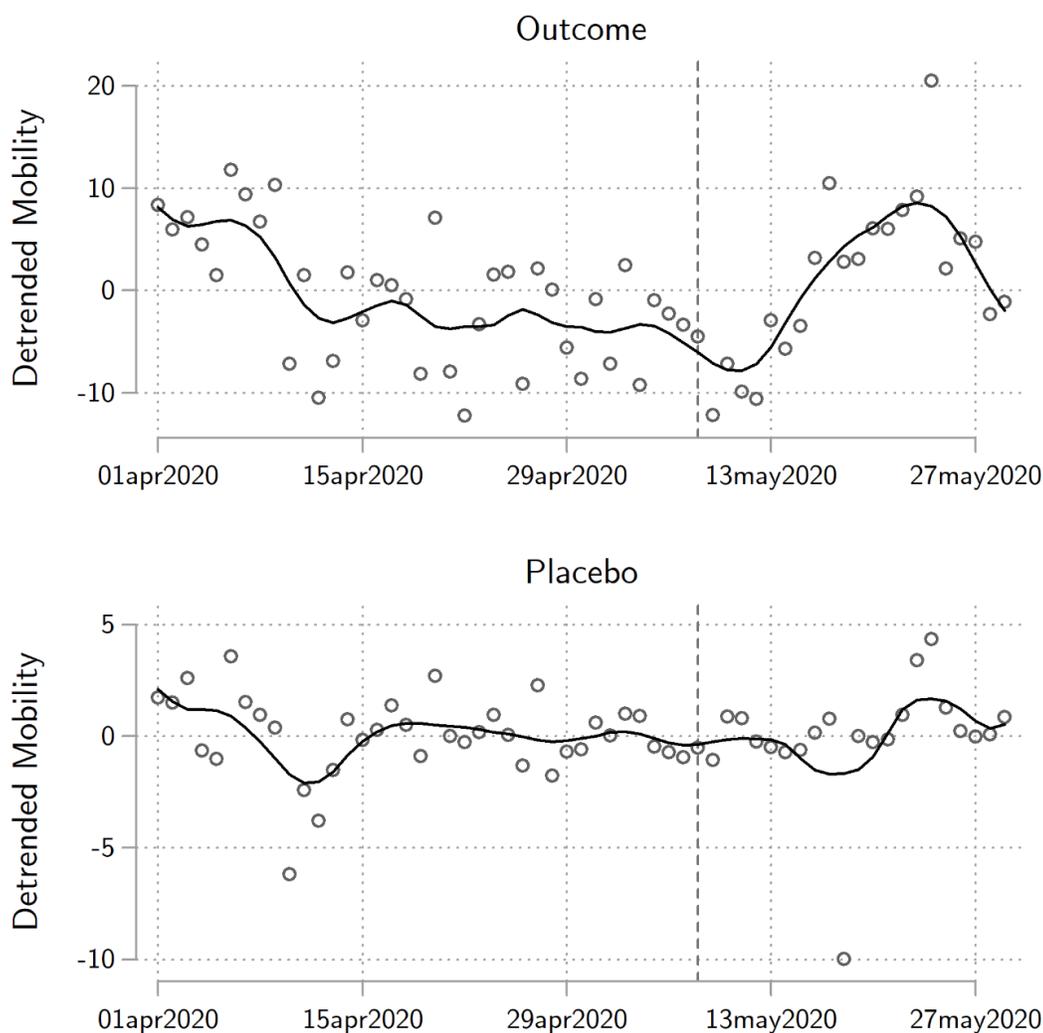


Figure 5. Averaged mobility data for retail, parks, and groceries (top panel); Averaged mobility data for workplaces, residential, and transit (bottom panel). Note: de-trended by extracting residuals from a model predicting mobility with a daily trend and fixed effects for day-of-the-week and month. Lowess smoothed at a bandwidth of 0.2.

⁴ Though transit mobility loads on both factors. We constructed an alternative index that includes transit. This measure is correlated with our own at 0.997.

We estimate lagged dependent variable models using OLS regression where we regress our outcome on its lag and a dummy variable were 1 indicates the period of the report release. We include a lag of the dependent variable to account for possible memory in mobility patterns – as in, mobility in one day will be systematically related to mobility in the next. We include fixed effects for month and day of the week. We control for average daily temperature and confirmed COVID-19 cases. The weather improved and newly confirmed cases fell during this period, so we want to account for these possible confounds which may encourage mobility. Weather data was taken from the Government of Canada’s historical data weather series for their station at Pearson Airport in Toronto. Daily confirmed COVID-19 cases was downloaded from the Government of Canada’s COVID-19 dashboard.⁵ Provincial governments began reopening their economies through this period as well. We add a dummy variable which takes a value of 1 after May 17 when Ontario and British Columbia re-opened their retail sectors. We cluster our standard errors by day of the week. We calculate the cumulative long run effect of our variables, which tells us the total effect of variable on the outcome over time ($X/(1-DV_{t-1})$).

We estimate a pair of models for our social distancing and placebo outcomes: one without a linear trend and another with the trend included. We see a trend towards more mobility throughout the period, which ultimately our explanatory variables aim to account for. The trend soaks up some of this effect. On the other hand, accounting for trending is necessary to meet the stationarity assumption of time series analysis and can help us control for unobserved, trending confounders. Consequently, we show our result both ways, where we expect the true estimate to be somewhere in between. Figure 2 plots our two series after soaking up trending, day of week, and month fixed effects. The dashed line represents the release of the report.

Google also offers province-level mobility series.⁶ Although the national-level jobs report is not expected to differentially effect citizens in various provinces, we can more effectively control for province-level temperature, daily COVID-19 cases, and re-opening policies.⁷ We use data from the weather station at the airport of the largest city in each province. We estimate time series cross sectional (TSCS) models with province-level fixed effects and panel-corrected standard errors. Including fixed effects with a lagged dependent variable leads to bias with small T. In the online supplement we provide evaluate the robustness of our finding with a static model using a Prais-Winsten correction for autocorrelation. Results are robust to this alternative specification.

⁵ Link: <https://www.canada.ca/en/public-health/services/diseases/2019-novel-coronavirus-infection.html>

⁶ Data is extremely sparse for Prince Edward Island, so this province is dropped from these analyses.

⁷ We use the opening date for retail as the indicator for province-level re-opening in these models. Note that Newfoundland and Nova Scotia did not re-open retail through this period, so they are scored ‘0’.

Table 1. Google mobility estimates, time series

	Social distancing		Placebo	
	1	2	3	4
Jobs report	6.96** (2.06)	4.39** (1.69)	0.35 (1.91)	-1.38 (2.04)
Re-opening	7.69** (2.53)	6.83** (2.38)	5.56 (3.27)	4.10 (3.31)
Mean temperature	0.65** (0.18)	0.55** (0.22)	-0.01 (0.19)	-0.09 (0.19)
Daily confirmed new cases	0.01* (0.00)	0.00 (0.01)	0.00 (0.00)	-0.00 (0.00)
DV _{t-1}	0.29** (0.11)	0.13 (0.16)	0.36*** (0.08)	0.17 (0.11)
Constant	-89.49***	-72.21**	-50.64***	-41.19***
Jobs report (CLRE)	9.79*** (3.46)	5.05* (2.61)	0.55 (3.00)	-1.66 (2.37)
Re-opening (CLRE)	10.80*** (3.11)	7.86** (3.20)	8.70* (5.16)	4.95 (4.05)
Mean temperature (CLRE)	0.91*** (0.35)	0.63** (0.32)	-0.02 (0.30)	-0.10 (0.22)
Confirmed daily cases (CLRE)	0.01 (0.01)	0.00 (0.01)	0.00 (0.00)	-0.00 (0.00)
Trend	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Day of the week FE	Yes	Yes	Yes	Yes
R ²	0.90	0.91	0.87	0.88
N	58	58	58	58

Note: Clustered standard errors in parentheses; * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$. CLRE = Cumulative long-run effect.

Google Mobility Results

The estimates are provided in Table 1. There is strong evidence that the report produced an increase in social distancing-related mobility.⁸ That is, the information shock is associated with increased mobility related to parks, retail, recreation, and grocery shopping, but not mobility related to residence, workplace, or transit. Model 1 shows that the jobs report period is associated with a 7 point increase in social distancing-related mobility ($p < 0.05$) after controlling for dynamics in re-

⁸ We ran a series of ten additional placebo tests by randomly drawing ten dates one week before or one week after the release of the jobs report. These serve as hypothetical report release dates with which we created identically constructed treatment measures. Across all ten tests we see no evidence of a treatment effect. We also include robustness tests by lagging the explanatory variables, and estimating static models. The estimates can be found in the left panel of Figure S2.

opening, COVID cases and temperature. The cumulative long-run effect is 10 points ($p < 0.01$). Re-opening is also associated with a cumulative increase in social distancing-related mobility of 11 points, while a one degree increase in average daily temperature is associated with a 0.91 cumulative increase in social distancing-related mobility ($p < 0.001$).

Table 2. Google mobility estimates, TSCS

	Social distancing		Placebo	
	1	2	3	4
Jobs report	12.75*** (3.01)	7.13** (3.22)	2.14 (1.51)	-0.07 (1.62)
Re-opening	0.00 (2.25)	-1.67 (2.31)	2.70*** (0.71)	2.18*** (0.70)
Mean temperature	1.49*** (0.18)	1.13*** (0.20)	0.24*** (0.06)	0.11** (0.05)
Daily confirmed new cases	-0.01 (0.01)	-0.01 (0.01)	0.00 (0.00)	0.00 (0.00)
DV _{t-1}	0.10* (0.06)	0.07 (0.06)	0.18** (0.08)	0.09 (0.08)
Constant	-72.22***	-12494.43***	-44.53***	-4908.55***
Jobs report (CLRE)	14.23*** (3.30)	7.63** (3.42)	2.60 (1.81)	-0.07 (1.78)
Re-opening (CLRE)	0.00 (2.52)	-1.79 (2.49)	3.29*** (0.86)	2.39*** (0.76)
Mean temperature (CLRE)	1.66*** (0.17)	1.21*** (0.19)	0.30*** (0.07)	0.12** (0.05)
Confirmed daily cases (CLRE)	-0.01 (0.01)	-0.01 (0.00)	0.00 (0.00)	0.00 (0.00)
Trend	No	Yes	No	Yes
Month FE	Yes	Yes	Yes	Yes
Day of the week FE	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes
R ²	0.64	0.66	0.79	0.81
Panels	9	9	9	9
T (mean)	57	57	57	57
N	510	510	516	516

Note: Panel-corrected standard errors in parentheses; * $p < 0.01$, ** $p < 0.05$, *** $p < 0.01$. CLRE = Cumulative long-run effect.

Model 2 adds the linear trend. Our effect sizes decrease. The report is associated with a 4.4 point increase in social distancing related mobility ($p < 0.05$), which rises to 5 points over time ($p < 0.1$). This is still a sizable 0.28 standard deviation increase in social distancing-related mobility after

accounting for trending, re-opening, temperature, and COVID caseload. Re-opening and temperature remain very important predictors of social distancing-related mobility. Re-opening is associated with 7.9 point cumulative effect on social distancing-related mobility ($p < 0.05$). A one degree increase in temperature is associated with a 0.6 point increase in social distancing related mobility over time ($p < 0.05$). The temperature rose from 1 to 24 degrees through May, meaning we would expect a 14 point increase in social distancing-related mobility as a result, or nearly a 0.78 standard deviation increase on this measure. In sharp contrast, there is no significant effect of the jobs report for our placebo test related to workplace, residential, and transit mobility.

Our TSCS estimates are very similar.⁹ The jobs report has a cumulative 12 point increase in social distancing related mobility ($p < 0.01$), which drops to 7 points when including the trend ($p < 0.05$). Province-level re-opening does not appear to have a significant effect on social distancing-related mobility in either model, but temperature remains an important predictor. A one degree increase in the temperature is associated with a cumulative 1.1 point increase in social distancing related mobility when including the trend ($p < 0.01$). In contrast, the jobs report has no significant effect on the placebo outcome in either model, though province-level reopening appears to matter. Irrespective of the powerful role of re-opening and temperature in shaping mobility through this period, we still identify a strong effect of the jobs report, in line with our individual-level findings.

Experiment 1

Our jobs report analyses convincingly show that social distancing dropped in Canada immediately following the jobs report, especially among the young. Some of this effect is likely direct – individuals increasingly perceive social distancing as costly and so changed their behavior accordingly. But some of this effect may be indirect. The information shock may have changed people’s expectations of others’ compliance and reduced their own compliance accordingly. We conduct a series of survey experiments to unpack this relationship.

We conducted our first experiment on an online sample of 2,495 Canadian citizens 18 years or older collected from March 25-30, 2020 by Dynata. Quotas were set by Canadian region (Atlantic, Quebec, Ontario, and West), age and gender, while weights are applied within region by gender and age as well for the following analyses. The survey for the first experiment was conducted in English only, and as a result Quebec is underrepresented in this sample. The demographic breakdowns for all three surveys can be found in the supplementary information.

Design

Experiment 1 was designed to evaluate the effect of prospective economic cost on expectations of social distancing compliance by other citizens. We exposed all of our respondents to a primer on social distancing found in the online supplement. Immediately following its passage, respondents were randomly assigned information about the economic consequences of social distancing. One group of respondents was told the economic consequences were uncertain:

⁹ Identical robustness tests for the TSCS models can be found in the right panel of Figure S2.

Economists, to date, are uncertain of the economic consequences of these changes to Canadian life.

Other respondents were informed of possible economic consequences, with groups receiving either information on job losses:

These changes to Canadian life are expected to have economic consequences.

Economists forecast a loss of [165/330/495/660],000 jobs by the third quarter of 2020.

Or the unemployment rate:

Economists forecast an increase in the unemployment rate to [6.7/7.7/8.7/9.7]% by the third quarter of 2020.

We randomly assigned both the scale of projected economic costs and the manner in which that information is communicated because we were uncertain as to the dosage required for a treatment effect and the most effective manner of conveying this information. We have strong priors that treatment effects will be higher as projected costs ratchet upwards and weak priors that people will be more responsive to job loss information rather than the unemployment rate, since some research has shown that people have little grasp of unemployment rates (Ansolabehere et al., 2011; Bullock et al., 2015).

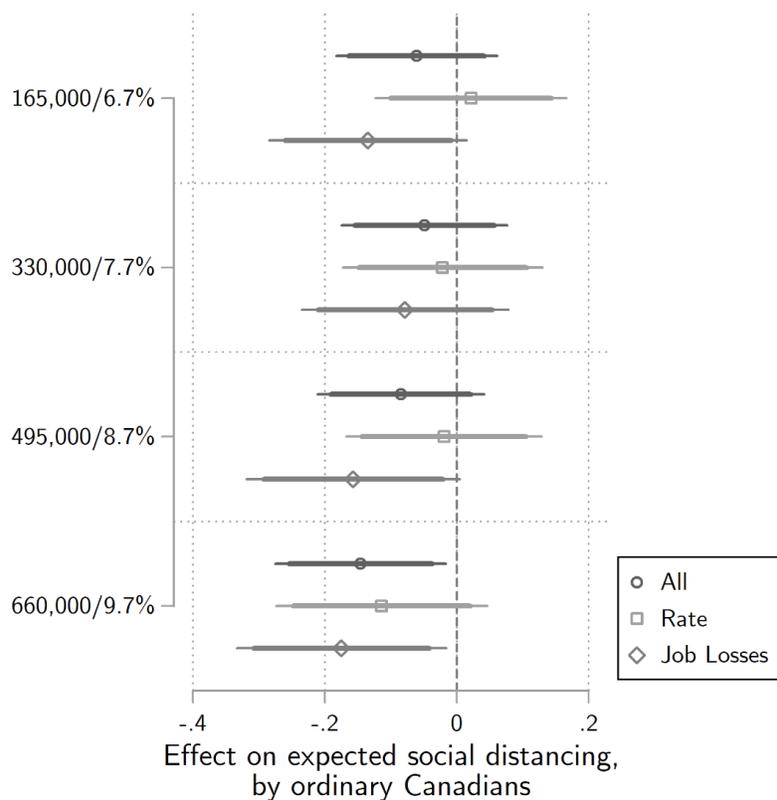


Figure 6. Effects of prospective economic cost treatment on expectations of social distancing by ordinary Canadians. Note: 95 and 90% confidence intervals.

Following the experiment, we asked our respondents how likely it is that “ordinary Canadians” would do the following actions for an additional two weeks even if they are not sick: 1) avoid public gatherings; 2) avoid in-person contact; 3) avoid restaurants, bars, and shops. These measures were scaled from not at all likely (0) to extremely likely (3). We construct an index from these items ranging from 0-1 (Cronbach’s alpha = 0.87). We standardize our outcome measure. We expect that respondents exposed to higher levels of projected economic cost will have lower expectations of social distancing compliance by other citizens.

Results

We plot the effects of our treatment conditions on social distancing expectations in Figure 6. We find that the highest cost condition, corresponding to a loss of 660,000 jobs or an unemployment rate of 9.7%, lowers expectations of social distancing compliance by 0.15 standard deviations ($p \sim 0.03$), consistent with H3. All other coefficients are signed in the correct direction, but are not statistically significant.

Our respondents appear to be comparatively more responsive to job losses. Information about a predicted loss of 165,000 jobs lowers expectations of the social distancing compliance by 0.13 standard deviations, though this is only marginally significant ($p \sim 0.08$). 495,000 jobs reduces the outcome by 0.16 standard deviations, which is also marginally significant ($p \sim 0.06$), while 660,000 jobs significantly reduces it by 0.17 standard deviations ($p \sim 0.03$). The corresponding effects for unemployment rate information are non-significant, though they are signed in the correct direction. The full estimates can be found in the online supplement.

Experiment 2

Negative economic information – at least as expressed as job losses – appears to lower people’s expectations that other citizens will comply with social distancing in the future (H3). It should also affect expectations of one’s own compliance as well (H4). We conducted a second experiment to provide a conceptual replication of our finding in Experiment 1 and to evaluate whether information regarding the prospective economic cost of social distancing lowers expectations of one’s own compliance, as is strongly suggested by our jobs report analyses.

We surveyed 2,499 Canadian citizens 18 years or older from April 2-6, 2020 using Dynata. Quotas were again set by Canadian region, language, age and gender, while weights are applied within region by gender and age as well. The survey was conducted in French and English, so the sample is representative along these dimensions as well.

Design

We randomly assigned our respondents information on the forecasted economic consequences of social distancing. The threshold for these costs was set based on the results of Experiment 1.¹⁰ All

¹⁰ Our expectation was that some of the treatment effect is mediated by changes in expectations of 1) other people’s compliance and 2) government policy. Consequently, we crossed our economic cost randomization with polling information about public social distancing compliance, and

respondents received a primer on social distancing that can be found in the online supplement. We then randomly assigned respondents into two conditions. One group received the following information on the projected economic costs of social distancing, while the other did not:

But these actions will come at a cost. Economists project a recession and for as many as 660,000 Canadians to be out of work by the end of the year. The economy will not recover until social distancing is relaxed.

Following the treatment, we asked our respondents how likely it is that they would do the following actions for an additional two months even if they are not sick: 1) avoid public gatherings; 2) avoid in-person contact; 3) avoid restaurants, bars, and shops; 4) avoid domestic and international travel. These measures were scaled from not at all likely (0) to extremely likely (3). We construct an index from these items (Cronbach's $\alpha = 0.85$).

We measure expectations of compliance by other citizens using the same question battery, but asking instead about "ordinary Canadians." We create an index from these measures (Cronbach's $\alpha = 0.84$). We also ask, on the same scale, how likely they think it is that federal and provincial governments would continue to adhere to and enforce current social distancing guidelines for two more months. Our outcome measures are standardized.

We also examine whether treatment effects for our primary outcome vary by age. Our expectation is that the treatment will have stronger effects on the self-expectations of social distancing compliance for younger respondents because they face higher costs relative to benefits from social distancing.

Effects could be heterogeneous across age for reasons that are unrelated to relative costs or benefits of social distancing. For example, we might expect anti-intellectual sentiment or ideology to moderate the treatment effect for reasons that have nothing to do with the costs and benefits of social distancing for individuals, and for there to be some correlation between these factors and age. News exposure may also moderate treatment effects because of the rapidly increasing coverage of the pandemic's economic effects, which may also be associated with age.

We present the results of a model interacting age with the treatment, as well as interactions between the treatment and controls for anti-intellectualism, left-right ideology, news exposure, education, income, location of residence, and region. Model estimates are provided in the online supplement.

government policy towards social distancing in order to manipulate these mediators. We provide a descriptive path model and the results of these experimental analyses in the online supplement. The results provide some suggestive evidence of mediation, but we were not successful in exogenously manipulating our mediators to allow for causal identification.

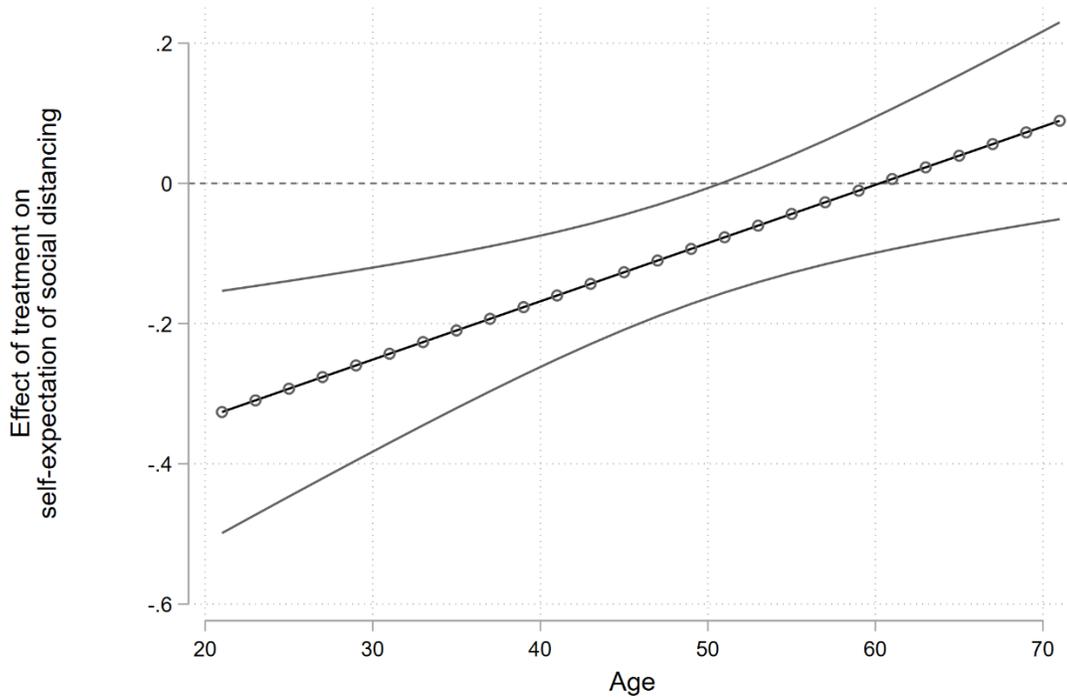


Figure 7. Heterogeneous effects of prospective economic cost treatment on self-expectations of social distancing across age. Note: 95% confidence interval, controlling for treatment effects across anti-intellectualism, left-right ideology, news exposure, income, urban/rural residence, education, and region.

Results

Once again, information about the high prospective economic costs of social distancing appears to affect expectations of compliance by other citizens. Respondents in the treatment condition scored 0.09 standard deviations lower in these expectations compared to those given no information ($p \sim 0.03$). The treatment also affected self-expectations of social distancing compliance. Respondents in the treatment condition reported self-expectations of compliance 0.1 standard deviations lower than those in the control ($p \sim 0.01$). In short, information about the high prospective economic cost of social distancing appears to lower expectations of social distancing compliance by other citizens and by oneself, echoing our April jobs report analyses. These results strongly support H3 and H4.

Figure 7 plots the estimated treatment effects across age controlling for potential confounders. Our prospective economic cost treatment is expected to lower self-expectations of social distancing by 0.29 standard deviations among 25 year-olds, while there is no detectable effect among those 53 and older. This interaction is statistically significant ($p \sim 0.003$). Information about the high prospective economic cost of social distancing lowers self-expectations of social distancing, primarily among younger respondents (H5), in line with our observational findings related to the April jobs report.

Experiment 3

Our second experiment provides a conceptual replication of Experiment 1, while also showing that this information reduces respondents' expectations of their own social distancing compliance. We believe that some of this latter effect is mediated by the former. We are only able to provide suggestive, non-causal evidence of mediation in Experiment 2 because we were unable to exogenously manipulate expectations of other people's social distancing (more details in the supplementary information). We cannot be sure whether or not expectations of other people's social distancing cause self-expectations of compliance rather than the reverse, nor can we rule out the presence of other, unmeasured mediators acting as confounders. So we conducted a third experiment to provide evidence of a causal arrow running from people's expectations of social distancing compliance by other citizens and expectations of their own compliance.

We surveyed 2,509 Canadian citizens 18 years or older between May 8-12, 2020 from Dynata. Quotas were again set by Canadian region, language, age and gender, while weights are applied within region by gender and age as well for the following analyses.

Design

We randomly assigned respondents polling information on the public's willingness to engage in various social distancing practices for an extended period of time to examine its effect on expectations of their own compliance.

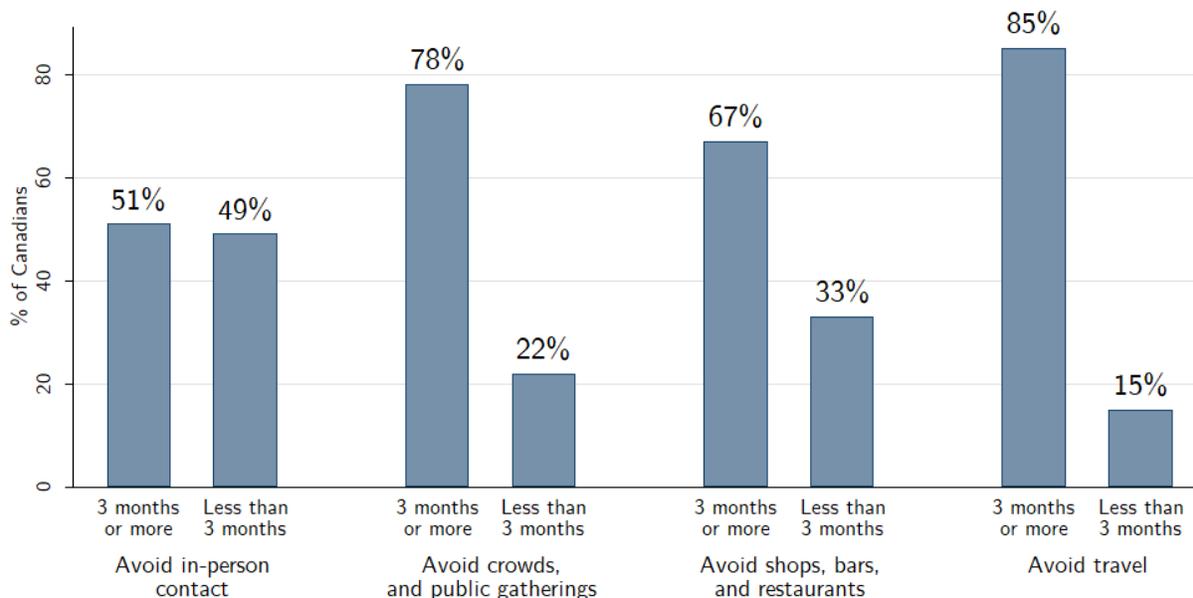


Figure 8. Treatment graph, Experiment 3

In a similar survey of 2,504 Canadians conducted from May 1-5, we asked respondents how many months they would be willing to: 1) avoid public gatherings; 2) avoid in-person contact; 3) avoid restaurants, bars, and shops; and 4) avoid domestic and international travel. We asked the

same about their expectations of “ordinary Canadians” (categories: less than one month, to more than six months).

Respondents underestimated the willingness of others to engage in long-term social distancing. 85% of respondents reported willingness to avoid travel for three or more months, but only 71% expected ordinary Canadians to do the same. 78% reported they would avoid public gatherings for the same period of time, but only 55% expected others to do the same, while 67% would avoid shops, bars, and restaurants, with only 43% expecting the same of others. Respondents were divided on their willingness to avoid in-person contact for three months or more, with only 51% signaling such an inclination. We exposed our survey respondents to a graph illustrating these results, shown in Figure 8.

We chose to expose respondents to information using the three month threshold to maximize the power of the treatment. Respondents generally have more accurate perceptions of public willingness to engage in social distancing for two months or more, while at 4 months there is substantial drop off in willingness to engage in social distancing (e.g. 34% are willing to avoid in-person contact, 48% are willing to avoid shops, bars, and restaurants). The three month threshold provides respondents with three clear cases of strong public support for long-term social distancing and one case – the avoidance of in-person contact – where the story is more mixed.

After the treatment we asked our respondents how many months they would be willing to: 1) avoid public gatherings; 2) avoid restaurants, bars, and shops; 3) avoid domestic and international travel; and 4) avoid in-person contact (categories: less than one month to over six months). The first three items were used to construct a scale for our primary outcome (Cronbach’s alpha = 0.83), while the in-person contact item was used for a placebo test. As a manipulation check, we asked the same of their expectations of “ordinary Canadians” and constructed similar measures (Cronbach’s alpha = 0.82). We standardize each our outcome measures.

Results

The polling information we provided increased the length of time respondents believed others would engage in social distancing. Exposure to the polling information is associated with a 0.21 standard deviation increase in expectations of social distancing compliance by other citizens ($p < 0.001$). This was not true of our placebo case – avoiding in-person contact – where polling indicated the public was divided ($p \sim 0.32$).

Exposure to the polling information is also associated with a 0.13 standard deviation increase in self-expectations of social distancing compliance ($p \sim 0.002$), consistent with H5. Again, this was not true for our placebo case. Respondents were no more willing to avoid in-person contact when given information that the public was divided on this question ($p \sim 0.32$). We use a path model to estimate the share of the treatment effect that is mediated by changes in respondent’s expectations of social distancing compliance by other citizens. Close to 90% of the effect travels indirectly through changes in expectations of other people’s compliance, as we would expect. The direct effect is not statistically significant.

Discussion

The COVID-19 pandemic is imposing catastrophic economic costs on citizens as they – in part voluntarily – give up personal liberties and forego social and economic activity to protect others. Citizens are being asked to engage in social distancing for an indeterminate amount of time as countries race toward developing a vaccine and effective, scalable testing and tracing mechanisms. We need to understand the degree to which this behavior is sustainable, and, if it is not, identify policy solutions.

We provide evidence that voluntary social distancing compliance is undermined by high economic costs. We leverage individual-level survey and aggregate-level mobility data in Canada before and after the release of the catastrophic jobs report on May 8, 2020. We show that self-reported social distancing declined at this point, while aggregate-level mobility rose (H1). We find this result to be the strongest among a set of respondents who face far higher costs relative to benefits of compliance – the young (H2).

Our experiments further illustrate that information about the high prospective economic costs of social distancing reduces both expectations of social distancing compliance by other citizens (H3) and by oneself (H4) – again, especially for the young (H5), echoing our findings related to the jobs report. We also provide causal evidence that expectations of social distancing compliance by others affects one's own expected social distancing compliance (H6). This suggests that high economic cost can affect self-expectations of compliance, in part, through changes in expectations of how others will behave.

The central implication of our findings is that social distancing is going to be hard to sustain as catastrophic economic costs mount. People may be less likely to engage in social distancing and less likely to believe others are compliant. Government policy needs to acknowledge the possibility that stringent, voluntary social distancing may not be sustainable and consequently accelerate investment in solutions that make society less reliant on such methods, such as testing and tracing. Alternatively, government can shift the imposition of costs to those who are mostly likely to receive benefits. This would imply stricter quarantines for at-risk populations.

Our analysis also reinforces the need of governments to do all in their power to cushion the economic fallout of the pandemic with fiscal and monetary policy, especially for those who face a disproportionate share of the costs of compliance with health directives, while reaping fewer of its benefits. The pandemic is not without distributional consequences and these citizens are more at risk of abandoning compliance in the future. Additional research should be conducted to examine whether government fiscal policy mitigates the effects we see here.

Our findings also signal the importance of expectations in reinforcing (or undermining) social distancing compliance. People tend to underestimate the willingness of others to engage in social distancing in the short and long-term, which may reduce their willingness to do the same. Informing citizens of this emerging norm provides a promising avenue of encouraging social distancing practices that are likely to come under increasing pressure as economic costs mount. Conversely, information that misleads citizens as to the behavior and sentiment of most citizens is unhelpful, such as escalating news coverage of extremely small and unrepresentative protests against government lockdowns.

Our analyses have some limitations. First, while we are confident that the April jobs report produced an exogenous shock of negative information, we cannot completely rule out the possibility that some other event occurred at the same time to produce a decline in social distancing compliance. Two factors mitigate against this concern: 1) we find equivalent effects when narrowing the bandwidth of our analysis to closer proximity to the release date; and 2) we still find effects controlling for COVID caseload and temperature changes in our aggregate-level analysis. Perhaps even more importantly, our experimental results reinforce our contention that negative economic information causes less social distancing.

Second, our experiments hinge on the reliability of measuring people's self-reported behavior, rather than observed compliance with social distancing in real time. Self-reports of behavior are prone to bias and random error in a number of different contexts, such as election turnout (Silver et al., 1986) and news exposure (Guess, 2015). We do, however, expect such error to be largely orthogonal to our experimental treatments and not a serious problem for the inferences we make here. Moreover, some recent work suggests that social distancing self-reports do a good job at predicting people's real world behavior (Gollwitzer et al., 2020).

A more serious problem is that we are measuring expectations of behavior, which is likely to deviate from their actual compliance in the future. We could imagine, for instance, that economic costs reduce expectations of social distancing but not actual behavior when it may put an individual at risk of infection. Ultimately, there is no way of overcoming this problem within the confines of a survey experiment. Our analyses related to the April jobs report that use both self-reported behavior and behavioral data helpfully supplement our experiments that focus on behavioral intentions.

Our findings highlight the critical importance of research maintaining societal compliance with social distancing guidelines and other recommendations of public health experts. The COVID-19 pandemic has placed enormous social and economic burdens on citizens. Much more research is needed on how citizens adapt to these circumstances and how government policy can ensure such adaptation is socially optimal. Such research can help us understand not only how social distancing compliance can be maintained, but how social dilemmas more generally can be solved.

References

- Akerlof, G.A., Kranton, R.E. (2000). Economics and identity. *Quarterly Journal of Economics*, 115, 715-753.
- Aldy, J. E., Kotchen, M. J., & Leiserowitz, A. A. (2012). Willingness to pay and political support for a US national clean energy standard. *Nature Climate Change*, 2(8), 596-599.
- Allcott, H., Boxwell, L., Conway, J. C., Gentzkow, M., Thaler, M., & Yang, D. Y. (2020). Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. NBER Working Paper No. 26946. DOI: 10.3386/w26946.
- Aklin, M., & Urpelainen, J. (2018). *Renewables: The politics of a global energy transition*. Cambridge, MA: MIT Press.
- Ansolabehere, S., & Konisky, D. M. (2014). *Cheap and clean: How Americans think about energy in the age of global warming*. Cambridge, MA: MIT Press.

- Ansolabehere, S., Meredith M., & Snowberg, E. (2011). Sociotropic voting and the media. In J. Aldrich & K. McGraw (Eds.), *Improving public opinion surveys: Interdisciplinary innovation and the American National Election Studies* (pp. 175-189). Princeton, NJ: Princeton University Press.
- Bauch, C. T., & Earn, D. J. (2004). Vaccination and the theory of games. *Proceedings of the National Academy of Sciences*, *101*(36), 13391–13394.
- Bechtel, M. M., & Scheve, K. F. (2013). Mass support for global climate agreements depends on institutional design. *Proceedings of the National Academy of Sciences*, *110*(32), 13763–13768.
- Bechtel, M., Genovese, F., & Scheve, K. (2019). Interests, norms and support for the provision of global public goods: the case of climate co-operation. *British Journal of Political Science*, *49*(4), 1333-1355.
- Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge: Cambridge University Press
- Bulle, R. J., Carmichael, J., & Jenkins, J. C. (2012). Shifting public opinion on climate change: An empirical assessment of factors influencing concern over climate change in the U.S., 2002–2010. *Climatic Change*, *114*(2), 169-188.
- Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, *10*, 519–578.
- Campbell, T. H., & Kay, A. C. (2014). Solution aversion: On the relation between ideology and motivated disbelief. *Journal of Personality and Social Psychology*, *107*(5), 809-824.
- Cornelson, K., & Miloucheva, B. (2020). Political polarization, social fragmentation, and cooperation during a pandemic. Available at: <https://www.economics.utoronto.ca/index.php/index/research/workingPaperDetails/663>
- Fenichel E. P. (2013). Economic considerations for social distancing and behavioral based policies during an epidemic. *Journal of Health Economics*, *32*(2), 440–451.
<https://doi.org/10.1016/j.jhealeco.2013.01.002>
- Fisman, D. N., & Laupland, K. B. (2009). The sounds of silence: Public goods, externalities, and the value of infectious disease control programs. *The Canadian Journal of Infectious Diseases & Medical Microbiology*, *20*(2), 39–41.
- Gollwitzer, A., Martel, C., Marshall, J., Höhs, J. M., & Bargh, J. A. (2020). Connecting self-reported social distancing to real-world behavior at the individual and U.S. state level. PsyArXiv Pre-print, 5 May. Available at: <https://psyarxiv.com/kvnwp/>.
- Guess, A. (2015). Measure for measure: An experimental test of online political media exposure. *Political Analysis*, *23*(1), 59-75.
- Kahan, D. M., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research*, *14*(2), 147-174.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*(3), 480-498.

- Gadarian, S. K., Goodman, S. W., & Pepinsky, T. B. (2020). Partisanship, health behavior, and policy attitudes in the early stages of the COVID-19 pandemic. SSRN Working Paper. Available at: <https://ssrn.com/abstract=3562796>
- Gaines, B. J., Kuklinski, J. H., and Quirk, P. J. (2007). The Logic of the Survey Experiment Reexamined. *Political Analysis*, 15(1), 1–20.
- Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., ... & Lesorogol, C. (2006). Costly punishment across human societies. *Science*, 312(5781), 1767-1770.
- Hopkins, D. J., Kim, E., & Kim, S. (2017). Does newspaper coverage influence or reflect public perceptions of the economy? *Research & Politics*, 4(4): 1-7.
- Jensen, U. T. (2020). Is self-reported social distancing susceptible to social desirability bias? Using the crosswise model to elicit sensitive behaviors. *Journal of Behavioral Public Administration*, 3(2). <https://doi.org/10.30636/jbpa.32.182>
- Leiserowitz, A. (2006). Climate change risk perception and policy preferences: The role of affect, imagery, and values. *Climatic Change*, 77, 45–72.
- Lewandowsky, S., & Oberauer, K. (2016). Motivated rejection of science. *Current Directions in Psychological Science*, 25(4), 217-222.
- Lodge, M., & Taber, C. S. (2013). *The rationalizing voter*. Cambridge: Cambridge University Press.
- Lowrey, A. (2020, April). Millennials don't stand a chance. *The Atlantic*. Available at: <https://www.theatlantic.com/ideas/archive/2020/04/millennials-are-new-lost-generation/609832/>
- Merkley, E. (2020). Anti-intellectualism, populism, and motivated resistance to expert consensus. *Public Opinion Quarterly*, 84, 24-48.
- Merkley, E., Bridgman, A., Loewen, P. J., Owen, T., Ruths, D., & Zhilin, O. (2020). A rare moment of cross-partisan consensus: Elite and public response to the COVID-19 pandemic in Canada. *Canadian Journal of Political Science*, 53(2), 311-318.
- Merkley, E., & Stecula, D. A. (2020). Party cues in the news: Democratic elites, Republican backlash and the dynamics of climate skepticism. *British Journal of Political Science*. Online Access. DOI: 10.1017/S0007123420000113.
- Motta, M. (2018). The dynamics and political implications of anti-intellectualism in the United States. *American Politics Research*, 46(3), 465-498.
- North, D. C. (1991). Institutions. *Journal of economic perspectives*, 5(1), 97-112
- Nowak, M., & Highfield, R. (2011). *Supercooperators: Altruism, evolution, and why we need each other to succeed*. Simon and Schuster.
- Ostrom, E. (1996). Crossing the great divide: coproduction, synergy and development. *World Development*, 24(6), 1073-1087.

- Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of economic perspectives*, 14(3), 137-158.
- Pasek, J. (2018). It's not my consensus: Motivated reasoning and the sources of scientific illiteracy. *Public Understanding of Science*, 27(7), 787-806.
- Pickup, M. A., Kimbrough, E. O., & de Rooij, E. A. (2020). Identity and the self-reinforcing effects of norm compliance. *Southern Economic Journal*, 86, 1222-1240.
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416), 427-430.
- Scruggs, L., & Benegal, S. (2012). Declining public concern about climate change: Can we blame the great recession? *Global Environmental Change*, 22(2), 505-515.
- Siegal, G., Siegal, N., & Bonnie, R. J. (2009). An account of collective actions in public health. *American Journal of Public Health*, 99(9), 1583-1587.
- Silver, B., Anderson, B., & Abramson, P. (1986). Who overreports voting? *The American Political Science Review*, 80(2), 613-624.
- Sinclair, B. (2012). *The social citizen*. Chicago: University of Chicago Press.
- Suhay, E. (2015). Explaining group influence: The role of identity and emotion in political conformity and polarization. *Political Behavior*, 37, 221-251.
- Tesler, M. (2018). Elite domination of public doubts about climate change (not evolution). *Political Communication*, 35(2), 306-326.
- Verity, R., Okell, L. C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N.,...Ferguson, N. M. (2020). Estimates of the severity of coronavirus disease 2019: A model-based analysis. *The Lancet Infectious Diseases*. Online Access.
- Young, L., & Soroka, S.N. (2012). Affective news: The automated coding of sentiment in political texts. *Political Communication*, 29, 205-231.

Supplementary Information

Contents

Lexicoder-based Sentiment Analysis.....	28
Survey Fielding Dates.....	28
Variable Descriptions.....	29
Robustness tests for Mobility Models.....	30
Sample Characteristics.....	31
Social Distancing Primers.....	31
Experiment 1 Estimates.....	32
Mediation Analysis – Experiment 2.....	32
Heterogeneous Effects Analysis – Experiment 2.....	36

Lexicoder-based Sentiment Analysis

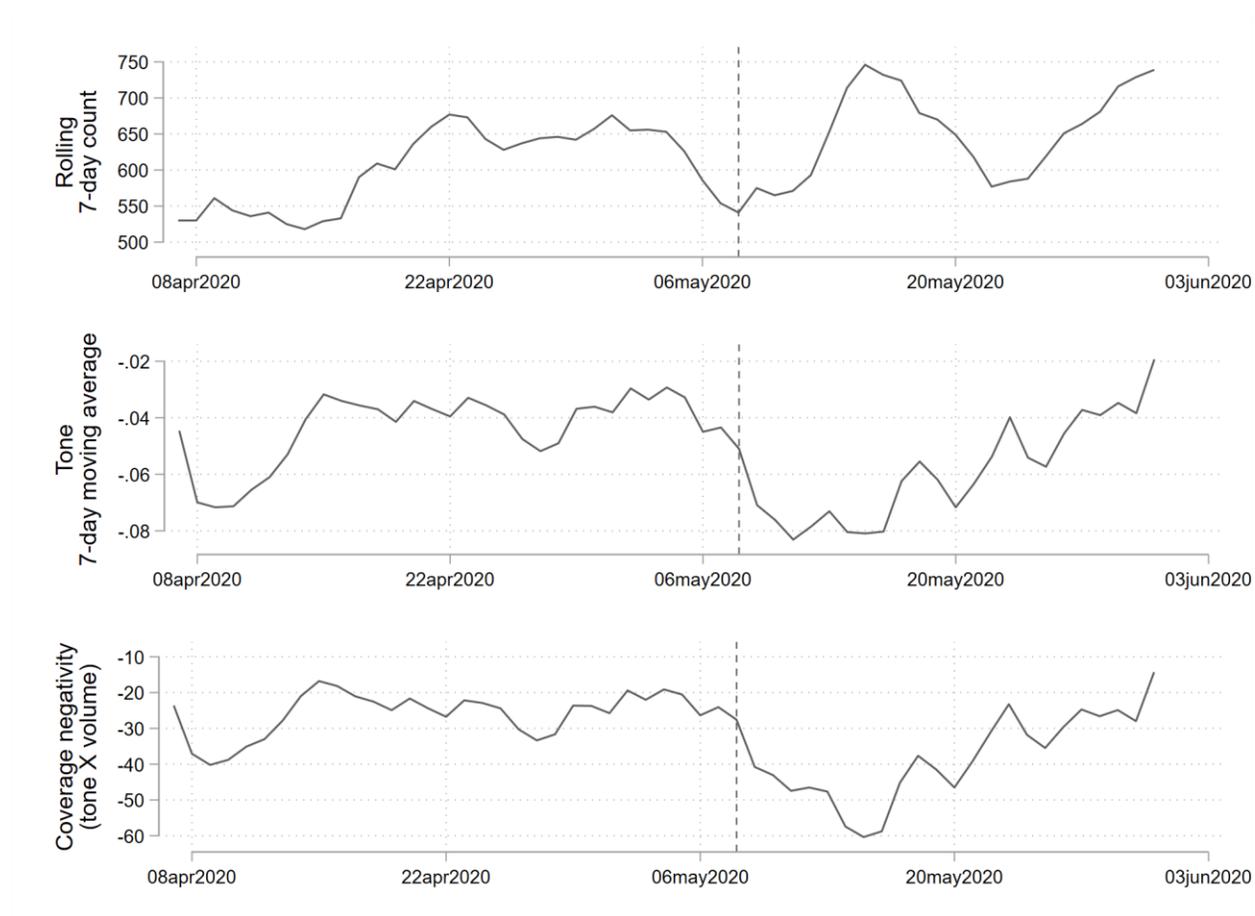


Figure S1. 7-day rolling sum of economic news content (top); 7-day moving average of headline sentiment (middle); composite coverage negativity index (bottom). Dashed line represents release of the April jobs report.

Survey Fielding Dates

Table S1. Survey waves and fielding dates

Wave	Fielding Date	N
1	April 2-6, 2020	2,489
2	April 9-11, 2020	2,493
3	April 16-19, 2020	2,489
4	April 24-29, 2020	2,515
5	May 1-5, 2020	2,512
6	May 8-12, 2020	2,514
7	May 21-27, 2020	2,527

Variable Descriptions

Table S2. Variable descriptions – observational analyses

Measure	Description
Social distancing	0-1; index of four behaviors over the past week: 1) Avoided bars, restaurants, and other places with crowds; 2) Avoided in-person contact with friends, family, and acquaintances; 3) Maintained 2 meters of distance from people as much as possible; 4) Avoided domestic travel.
Social distancing – alternative	0-1; first factor of a PCA: 1) worked from home; 2) avoided grocery stores at peak times; 3) stocked up on provisions; 4) switched to virtual meetings; 5) switched to online shopping; 6) avoided public transit; 7) Avoided bars, restaurants, and other places with crowds; 8) Avoided in-person contact with friends, family, and acquaintances; 9) Maintained 2 meters of distance from people as much as possible; 10) Avoided domestic travel.
COVID-19 risk perceptions	0-1; How serious of a threat do you think the coronavirus is to Canadians? (very, somewhat, not very, not at all)
Anti-intellectualism	0-1; Below is a list of groups and institutions in society. Please tell us the degree to which you trust or distrust members of these groups or institutions: 1) Experts; 2) Economists; 3) Scientists; 4) Doctors and medical professionals; 5) University professors (distrust a lot, distrust somewhat, neither, trust somewhat, trust a lot, don't know)
Ideology	0-1; 1) The government should take measures to reduce differences in income levels; 2) Protecting the environment is more important than creating jobs; 3) Canada should increase the number of immigrants it admits each year; 4) People who don't get ahead should blame themselves, not the system; 5) The government should see to it that everyone has a decent standard of living (Strongly, somewhat, neither agree/disagree). Each item coded in left-wing (-1) and right-wing (1) direction. Don't knows and neither coded as neutral (0)
News exposure	Logged sum of exposure to following outlets in past week: 1) CBC; 2) CTV; 3) Global; 4) CityNews; 5) Globe and Mail; 6) National Post; 7) Toronto Star; 8) Local newspaper; 9) TVA (French-only); 10) TV5 (French-only); 11) La Presse (French-only); 12) Journal de Montreal (French-only); 13) Journal de Quebec (French-only); 14) Le Devoir (French-only); 15) Radio-Canada (French-only); 16) Rebel Media; 17) National Observer; 18) Toronto Sun; 19) The Tyee; 20) Post Millennial; 21) APTN; 22) True North News; 23) Press Progress; 24) Huffington Post; 25) Other

Education	Highest level of education: no schooling; some elementary school; completed elementary school; some secondary/high school; completed secondary/high school; Some technical, community college, CEGEP, College Classique; Completed technical, community college, CEGEP, College Classique; Some university; bachelor's degree; master's degree; professional degree or doctorate; don't know
Age	Age in years
Urban/rural	Thinking about the place where you live, what word best describes it: A large city, a medium sized city, a large town, a small town, a rural place.
Gender	1= female
Region	Province of residence: Atlantic = Newfoundland and Labrador, Prince Edward Island, Nova Scotia, New Brunswick; Quebec; Ontario; West = Manitoba, Saskatchewan, Alberta, British Columbia

Robustness tests for Mobility Models

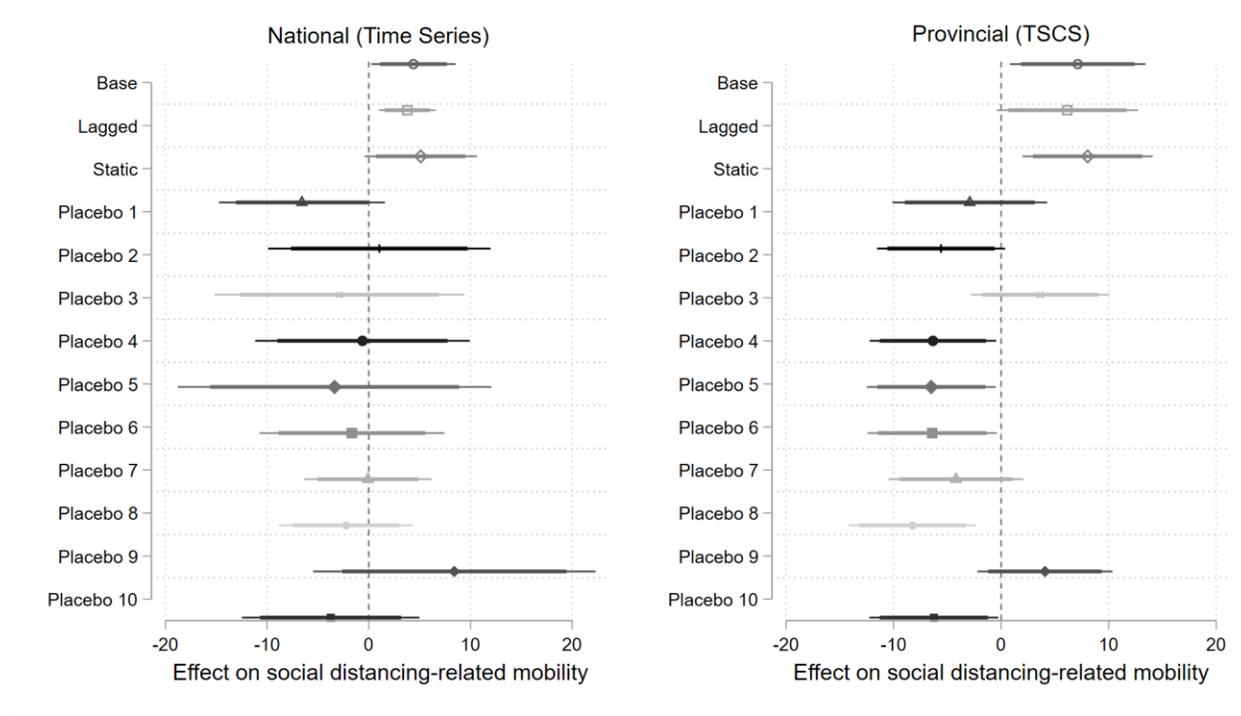


Figure S2. Robustness tests for national-level (left) and provincial-level mobility models (right). Each model controlling for temperature, daily COVID-19 cases, reopening, lagged social distancing-related mobility, month and day-of-week fixed effects, and a trend. National models have standard errors clustered by day-of-week, provincial models have panel corrected standard errors.

Sample Characteristics

Table S3. Sample characteristics

		1	2	3
		March 25-30	April 2-6	May 8-12
Female		51.5	51.1	51.5
Age	18-34	27.3	27.2	25.2
	35-54	34.2	33.5	34.2
	55+	38.6	39.3	40.5
University educated		40.0	36.2	38.8
French		0.0	20.4	20.3
Region	Atlantic	7.9	6.9	6.9
	Quebec	10.3	22.9	22.7
	Ontario	45.1	38.3	38.6
	West	36.8	31.9	31.7
2019 Liberal vote		39.4	35.8	37.1
N		2495	2499	2509

Social Distancing Primers

Experiment 1

The coronavirus or COVID-19 is highly contagious. A person who carries the virus spreads it to an average of 3 people.

In an effort to control the spread of the virus, public health officials have called on Canadians to engage in social distancing where they limit in-person contact with high risk individuals, avoid common greetings, such as handshakes, steer clear of crowded places and non-essential gatherings, and keep a distance of two meters from others as much as possible.

Experts believe that these efforts can protect Canadians by limiting the coronavirus' strain on the health care system.

Experiment 2

In an effort to control the spread of the coronavirus, public health officials have called on Canadians to engage in social distancing where they keep a distance of two meters from others as much as possible, limit in-person contact with high risk individuals, avoid travel, and steer clear of crowded places like bars, shops, restaurants, and non-essential gatherings. This is best achieved by staying at home as much as possible.

Experiment 1 Estimates

Table S4. Experiment 1 estimates

	Both		Unemployment Rate		Job Losses	
	Coef.	SE	Coef.	SE	Coef.	SE
165,000/6.7%	-0.06	0.06	0.02	0.07	-0.13*	0.08
330,000/7.7%	-0.05	0.06	-0.02	0.08	-0.08	0.08
495,000/8.7%	-0.08	0.06	-0.02	0.08	-0.16*	0.08
660,000/9.7%	-0.15**	0.07	-0.11	0.08	-0.17**	0.08
Constant	0.07	0.05	0.07	0.05	0.07	0.05
N	2495		1485		1502	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Mediation Analysis – Experiment 2

The results of Experiment 2 show that information regarding the high prospective economic cost of social distancing lowers both expectations of compliance by other citizens and expectations of their own compliance. We expect that at least some of the treatment’s effect on self-expectations of social distancing compliance is mediated by changes in expectations of other people’s behavior. We attempted to test this theory by using a 2x2x2 between subjects, manipulation of the mediator design in Experiment 2.

In addition to assigning information about economic cost, we also randomly assign information signaling the public’s support for social distancing using data from our surveys:

Upwards of 85% of Canadians signal continued willingness to engage in social distancing practices according to public opinion polls.

Our expectation is that the negative effect of economic cost will be weaker among respondents assigned to this condition since we have set the mediator at a high level.

Mediation analysis poses extraordinary challenges. Experimentally manipulating the mediator is not a panacea because its manipulation could activate other traits or considerations outside of your causal logic that influences the outcome (Bullock et al., 2010). In this case, we have reason to believe that high prospective economic costs and lowered expectations of compliance by other citizens could create an expectation that governments will fail to continue to enforce social distancing guidelines. This, in turn, could influence prospective social distancing for reasons that have nothing to do with the failure of others to contribute to public goods through social distancing. Consequently, we randomly assign information that federal and provincial government intend to stand by social distancing as long as public health officials advise, while those in the control condition received no such information:

Federal and provincial governments intend to stand by social distancing guidelines as long as public health officials deem them necessary, which may be for weeks, if not months.

Our expectation is that the effect of economic cost will be weaker among respondents given this information. It should also directly affect expectations of one's own social distancing.

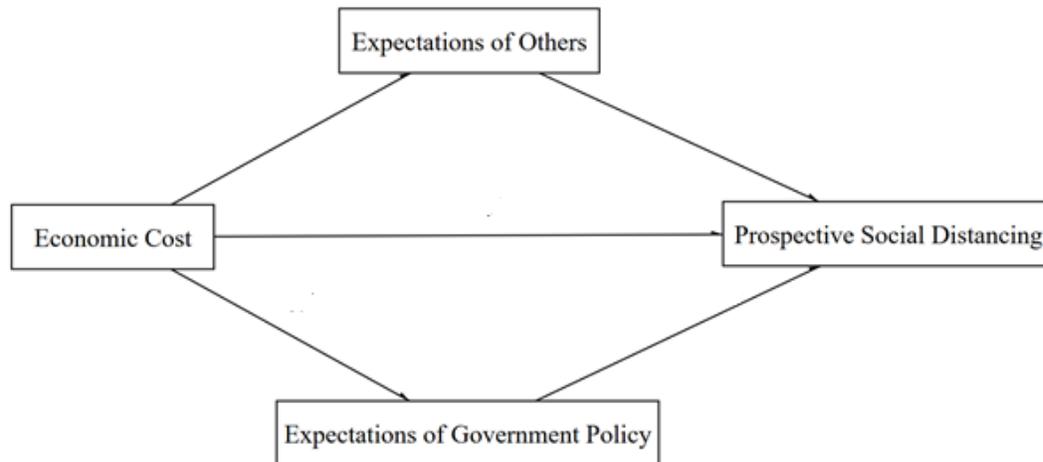


Figure S3. Path model

Method

We take two approaches to studying mediation in Experiment 2. First, we estimate the path model shown in Figure S3 that specifies a mediating effect of expectations on self-expectations of social distancing compliance. Path models allow for the testing of directional effects specified by a particular theory. In this case, it will allow us to estimate the direct effect of our exogenous economic cost treatment on self-expectations social distancing, and the indirect effect that flows through changes in expectations of other citizens' behavior and of government policy.

This analysis, however, is intended only to be descriptive. Path models require strong assumptions about causal direction and an absence of confounding that is nearly impossible to meet in social science (Bullock et al., 2010). We cannot rule out the existence of unmeasured mediators acting as confounds by being correlated with our mediators and the outcome. Nor can we be sure that the causal direction is from expectations of compliance by other citizens to self-expectations of compliance rather than the other way around. In short, we have causal identification between our economic cost treatment and our mediators, and between our treatment and our outcome, but not between our mediators and the outcome.

As a result, we randomly assigned information to respondents to manipulate our mediating variables. We estimate the following models interacting each mediator treatment with economic cost. Our expectation is that treatment effects would be weaker when providing information that sets the mediator at a high level for respondents. A positive and significant coefficient on β_3 and a negative coefficient on β_1 would provide support for a mediating relationship:

self – expectations

$$= \alpha + \beta_1 \text{cost} + \beta_2 \text{public support} + \beta_3 \text{cost} * \text{public support} + \beta_4 \text{government support} + \beta_5 \text{cost} * \text{government support} + \varepsilon$$

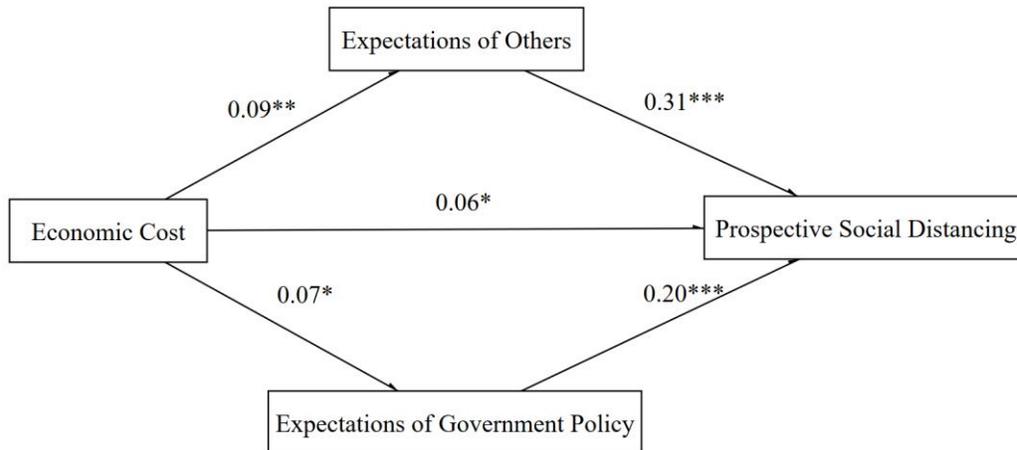


Figure S4. Direct effects of path model

Results

The direct effects of our path model are shown in Figure S4. The results indicate that about 60% of the total effect flows directly from economic cost to prospective social distancing (0.06 standard deviations, $p \sim 0.09$), while 40% flows indirectly through expectations of others and government policy (0.04, $p \sim 0.02$). That is, these results suggest that expectations of public and government compliance with social distancing partially mediate the effect of high prospective economic cost on self-expectations of social distancing.

This analysis, however, is descriptive. It does not provide causal identification between the mediating variables and the outcome. Unfortunately, our experimental results provide no evidence that treatment effects were weaker among respondents given information related to government policy or public support for social distancing, as is shown in model 1 of Table S3. In fact, the signs on the interaction terms are in the wrong direction.

This is likely a result of our treatments failing to manipulate our mediators of interest. Model 2 shows that information on public support for social distancing failed to increase expectations of social distancing compliance by ordinary citizens ($p \sim 0.24$). Model 3 illustrates that expectations of government policy are similarly uninfluenced by information that government social distancing guidelines will continue indefinitely in response to advice of health experts ($p \sim 0.63$). Both of the estimates are signed in the wrong direction. We can thus only say that our analyses provide suggestive evidence of mediation, but we do not have causal identification of the link between our mediator and the outcome. Therefore, we turn to Experiment 3 for evidence of a causal link between expectations of compliance by other citizens and self-expectations of compliance.

Table S5. Experimental mediation analysis and manipulation checks

	Expectations		
	Self 1	Others 2	Policy 3
Economic Cost Treatment	-0.02 (0.07)		
Government Support Treatment	0.06 (0.06)	-0.02 (0.04)	
Cost * Government Support Treatment	-0.11 (0.08)		
Public Support Treatment	-0.02 (0.06)		-0.05 (0.04)
Cost * Public Support Treatment	-0.05 (0.08)		
Constant	0.03	0.02	0.01
N	2499	2499	2499

Standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Heterogeneous Effects Analysis – Experiment 2

Table S6. Model estimates, Experiment 2

	Controls	
	Coef.	SE
Cost	-0.118	0.29
Age	0.007***	0.002
Cost * Age	0.008***	0.003
Anti-intellectualism	-0.130***	0.038
Cost * Anti-intellectualism	-0.114**	-0.112*
Ideology	-0.030**	0.013
Cost * Ideology	0.006	0.02
News Exposure	-0.012	0.055
Cost * News Exposure	0.075	0.081
Education	0.042**	0.017
Cost * Education	-0.017	0.024
Income	0.017	0.021
Cost * Income	-0.024	0.028
Urban/Rural	-0.024	0.022
Cost * Urban/Rural	0.072**	0.035
Quebec	0.245**	0.114
Ontario	0.126	0.112
West	0.114	0.113
Cost * Quebec	-0.400**	0.162
Cost * Ontario	-0.350**	0.162
Cost * West	-0.418**	0.163
Constant	-0.582***	0.208
N	2300	

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$