Why can governments use regulation but not technology?

by Peter John Loewen, John McAndrews, Jad El Tal, Kulsoom Khalid, Haobo Chen, Nick Thompson, Anthony Moniuszko, Erika Porco, Mitchell Anderson, and Sapphira Thompson-Bled

A pre-read paper from the Schwartz Reisman Institute for Technology and Society in partnership with the Munk School of Global Affairs and Public Policy

Innovating AI Governance: Shaping the Agenda for a Responsible Future
December 2020

Introduction

The response of governments and citizens to the COVID-19 pandemic has been a demonstration in the power of governments to regulate and the ability of citizens to coordinate for massive collective action. But it has also demonstrated the failure of either group - governments or citizens - to use technology towards shared ends.

The COVID-19 pandemic has brought governments back to the forefront of Western societies. On the one hand, the power of governments has been demonstrated through their capacity to impose lockdowns and widespread restrictions on civil liberties. The scale and scope of these lockdowns is remarkable. Governments have imposed severe restrictions on movement, assembly, and commerce. They have in some cases even denied citizens access to often critical medical services. Such lockdowns have not been limited to a small number of countries, and certainly not only to non-democratic countries. Instead, across the democratic world, states have imposed (and sometimes reimposed) severe restrictions on the civil liberties of citizens.

On the other side of the ledger, governments have demonstrated their ability to act as fiscal and economic backstops, effectively staving off a global recession while confronting a global pandemic. For example, the French government has offered immediate fiscal stimuli eight different times since March, totaling €110 billion (source). Similarly, the federal German government created an Economic Stabilisation Fund (WSF) worth 8.3% of its GDP in immediate fiscal stimuli, totalling approximately €290 billion. Even more fiscally conservative governments such as the United Kingdom adopted job retention schemes, small business grants and welfare measures totalling £176.7 billion (same source as above). In Canada, the federal government is on track to hitting the largest budgetary deficit at C$328.5 billion for the 2020/21 fiscal year, which includes C$226 billion in pandemic relief spending (source). In comparison, last year's budgetary deficit was around C$20 billion. The economic slowdown that has necessitated these backstops is itself something to behold. As a result of lockdowns, dampened consumption, and worldwide slowdown, the United Kingdom saw its unemployment rate climb from 3.9% to 4.5%. In the United States, the jump was from 3.5% to 14.7% (source). In Canada, it increased from 5.6% to 13.7% (Source). Importantly, the vast majority of this decline was attributable to government choices, not to market forces.

This has been an impressive show of force by governments, and it has had positive effects. Projections from public health officials - across Western countries - painted a grim picture of the effects of an unchecked SARS-CoV-2 virus. Lockdowns and other measures seem to have prevented larger fatality numbers. What is more, in countries that have taken laxer approaches to restricting activity and social contact - for example, Sweden and several US states - is that there have been noticeably higher numbers of deaths, and no effective movement towards herd immunity.

All of this action has been taken in the face of persistent and often growing beliefs in the ineffectiveness of governments and a concomitant decline in trust within governments. Once again, governments appear central to the action. But governments have used blunt tools. Complete shutdowns of economies and the suspension of civil liberties are massive costs to pay, both economically and normatively. These have been done in the service of the simple outcome of limiting contact between people.

To be sure, the coronavirus is a challenging virus from a public health perspective, because it combines both a relatively high rate of infectiousness (though not as high, for example, as the measles) with a relatively high fatality rate (though not as high, for example, as Ebola) (see Christakis 2020 for a nice discussion of this combination of factors). What is more, is that since the virus can spread through asymptomatic and presymptomatic individuals, it can be carried unknowingly. And because it can spread via aerosols, it may not be sufficient to simply ensure spacing between people indoors. To prevent the spread of the disease (especially community transmission), governments have had to stop people from being close to one another. But this has been massively costly.

Just as governments have taken extensive measures, citizens have also taken extensive measures. While we demonstrate this more comprehensively in the next section, an overview of the response of citizens is impressive. Across the Western world, citizens have recognized the severity of the COVID-19 pandemic. For example, YouGov data has found that at least 40% of people in the Western world have consistently been “very” or “somewhat” scared that they will contract COVID-19 since March, and the trend is increasing as we enter a second wave. It may seem unremarkable that such consensus by citizens exists, but when put in relief against widespread disbelief about climate change and more generally anti-intellectualism and opinion polarization, the importance of this widespread consensus on concern over COVID-19 is notable.

Citizens’ behaviours have matched their beliefs. The same data has shown that more than 60% of people have avoided crowded places and improved personal hygiene (such as frequent handwashing). This compliance by citizens often comes at a very substantial economic cost through lost employment, not to mention the deep social and psychological costs of living at a distance from family members, friends, and coworkers.

And, finally, citizens have rewarded governments for taking bold actions. In Canada, for example, the first three governments to contest elections in the pandemic have won greater re-election victories. More broadly, data show that governments have experienced sustained increases in citizen approval, perhaps especially immediately following the exercise of lockdown actions (Bol et al 2020).

If governments were impressive in locking down their economies, they were also impressive in opening them back up. Across the Western world, governments have shown an ability to open up economies, in a fashion differentiated across sectors, and, through the summer at least, to keep infection rates low. Citizens have similarly shown a capacity to respond to new public health guidelines, adopting masks when recommended, tightening and expanding social circles as guided, and taking other preventative measures at the same time as gradually returning to more normal life.

All of this then suggests something remarkable: in the maws of a major crisis, citizens are looking to governments (Teichmann et al 2020) and are consenting to big, involved, active governments. And governments are responding in kind. For those who have long argued about the centrality and necessity of governments, this is perhaps a time of great success.

From another perspective, however, this represents a major failure of governments. Governments have relied on massive and blunt instruments – like total shutdowns accompanied by massive economic support – because, at least in part, they cannot adopt smarter, technology-driven solutions. In particular, in many countries, testing has taken far too long to come online, and the testing that has come online, is often insufficiently precise and/or rapid to obviate the need for more inefficient practices of social distancing under the uncertainty of infection. Many governments have been reluctant to use technologies to enforce quarantines on infected individuals (and perhaps those within one degree of contact with them). And, most notably, governments have been very ineffective at driving sufficiently wide scale adoption of contact tracing applications to allow them to play a central role in controlling the pandemic, despite the potential of these tools to help combat the virus (Ferretti et al 2020).

Why is it that governments, which have otherwise shown both great effort and great effectiveness in the use of other tools, have not been able to drive widespread adoption of apps which would help contain the virus? (A related question, which we do not explore here, is why they would also choose arguably more ineffective but privacy preserving Bluetooth proximity tracing applications, as opposed to using the location tracing capacities available in smartphones). The purpose of this paper is to explore this in more depth. We do this in four stages. In the next section, we present unique Canadian data on citizens' reactions to COVID, as well as data on government policies from across the OECD. After this introduction, we briefly expand the review of both citizen and government action. This leads up to a consideration of the use of contact tracing apps to combat the spread of the coronavirus. As we show, there has been limited uptake of these apps by citizens. And, to the degree that governments have been able to encourage widespread usage of the apps, it has not been sufficient to trade-off economic regulation for widespread app usage. In the section following this, we take a deeper dive into understanding why some citizens in Canada have been willing to use a contact tracing app while others have not. Then, focusing on the non-adopters, we

explore potential interventions or encouragements which might increase adoption. By looking at which interventions work, we can go some distance in understanding what governments need to do to convince a much larger degree of uptake among citizens. This is essentially, then, a case study in the unwillingness of citizens to partner with the government in the use of an app. The paper then concludes with three hypotheses or intuitions about the use of technology versus the use of cruder, more expensive, and blunter instruments of policy.

Data

The data we rely on in this paper, comes from two related projects, both run out of PEARL (the Policy, Elections, and Representation Lab), and funded by the Schwartz-Reisman Institute for Technology and Society and the Munk School of Global Affairs & Public Policy. The first is Canadian public opinion and behavioural data, collected as a part of the Media Ecosystem Observatory, a joint project with two McGill labs, exploring the relationship between information, media, and citizen responses to the COVID-19 pandemic. This project runs a weekly online survey on a representative sample of Canadians, generally varying in size from 1500 to 2500 respondents. Our surveys contain a core module of questions related to beliefs and concerns about COVID, objective health data, social distancing behaviours, etc. A large number of experimental modules are swapped in and out depending on the week. Running since March 2020, we have now conducted 27 waves of data. The modelled data we present below were collected in the 25, 26, and 27 waves (running from October 8 to November 2, 2020). The social distancing and masking data we present below come from all waves of the survey. Our data are weighted using iterated proportional fitting, conditional using census observables on age, gender, and region.

Our government policy response data comes mainly from a dashboard that is run by PEARL. This dashboard measures the openness of economies on nine different dimensions in 52 different jurisdictions: 34 OECD countries, 5 US states, 10 Canadian provinces and 3 Canadian territories. These measurements are reported weekly in a qualitative summary of actions that each jurisdiction is taking and what they are considering as they reopen their economies and manage COVID-19. Since May 22, we have quantitatively classified how open each jurisdiction is based on nine dimensions: Stores/Non-Essential Businesses, Schooling/Youth Activities, Eateries, Leisure Activities, Contact Services, Manufacturing/Construction, Gatherings, Borders/Movement, and Cultural. Each sector is assigned a level of reopening from Level 1 (most restrictive) to Level 5 (least restrictive).

In order to have a clearer understanding of trends across certain regions, we have grouped certain jurisdictions based on their OECD designations (Northern Europe, Central Europe, etc.). We have chosen to highlight the 5 main states that have strong economic, people-to-people, and industry ties to Canada. We decided to evaluate how each province and territory is reopening in Canada because of the epidemiological divergence between regions (the

"Atlantic Bubble"). In May, we realized that jurisdictions were reopening their economies incrementally based on certain sectors. Accordingly, we decided to codify government policies across the nine aforementioned dimensions. Generally, Borders/Movement continue to be the most closed dimension throughout the last six months, and Manufacturing/Construction is the most open. While jurisdictions were quick to close schools in March, governments are now more hesitant to close schools in a bid to keep the economy as open as possible. For example, France recently reinstated a nationwide lockdown to tackle the second wave. Non-essential businesses have closed their doors and employees work from home where possible, yet going to in-person classes is still mandatory for grade school students.

Every week, policy analysts in the Lab sift through news articles and government press releases across all jurisdictions. They first describe the epidemiological situation based on each jurisdiction's metric of analysis (R-naught, 7-day average of cases, number of hospitalizations, positivity rate of tests etc.). Then, they evaluate whether a certain dimension has become more open, less open or stayed the same, based on an evaluation criteria for each dimension. Finally, a jurisdiction's level of reopening (scale from 1.0 to 5.0) is calculated based on the average of reopening across all dimensions for each jurisdiction. Detailed data and qualitative summaries are available at www.reopeningaftercovid.com.

Policy and Citizen Responses

In the introduction, we sketched a picture of three broad empirical facts about government and citizen responses to the pandemic, thus far. First, governments have been both willing and able to regulate large domains of economic and social activity to try to constrain the spread of the coronavirus. Second, citizens have similarly been willing to undertake personally costly actions. Third, despite this, governments have had little success in convincing citizens to take app contact tracing applications, and certainly not at a level which would allow those apps to generate a trade off between adoption and less regulation. In this section, we substantiate these empirical claims more concretely.

Figure 1 demonstrates compliance with social distancing and masking among Canadian citizens, from early spring to the present. Three findings are of note. First, citizens were, from the beginning, willing to comply with social distancing. Our measure relies on four social distancing actions, and so a score of 50%, for example, indicates that the average citizen reports complying with two of four actions. From the beginning of our data collection in March, we have found evidence of citizens being largely compliant, and this has continued to the present day, even when some of the actions may prove less necessary given our advancing knowledge of the transmission of the virus. Second, when new recommendations came online, citizens largely compiled with them, as shown by the sharp increase in masking after this become an official recommendation of public health officials. Third, this compliance has continued to today. Canadians, it appears, have been doing their part.

Figure 2 demonstrates the actions governments in Canada have taken in response to COVID. Here, our data begin in early June, when case counts were once again under control. As we can see, then, governments moved over the summer to relax restrictions, and most provinces converged on relatively open economic activity. However, variance still remained between provinces, sometimes substantially. And now, as case counts again rise, there is an almost monotonic shift towards more closed policies.

The pattern is even more dramatic when we consider all OECD countries. Across the developed world, governments showed a remarkable ability to rapidly change their closure or opening policies. And, as in Canada, there has again been a sharp downward movement.

None of this - widespread citizen compliance or regulatory agility from governments should be regarded as a small matter. On the citizen side, governments regularly face difficulties in convincing citizens to engage in voluntary actions, even on a small scale. And on the regulatory side, the increasing complexity and diversity of the administrative state normally renders rapid changes to policies on commerce and labour impossible. And yet, across the OECD, governments have been rapidly changing and enforcing policies.

This is all true and impressive, except in the domain of technology, especially contact tracing applications. As Figure 4 shows, as of November 2020, of all the countries in our sample, only three of 24 countries have download rates greater than 40%. Fully half have download rates below 20%. And the maximum rate, in New Zealand, does not reach a majority of the population. Importantly, contact tracing apps require a certain minimum rate of adoption to be useful in combatting the spread of the coronavirus, and no country has reached that critical diffusion rate, estimated to be 60-70% of a population (Hinch et al 2020). Indeed, even in countries with high uptake rates, the efficacy of apps has proved limited (Johnson 2020), suggesting for some that older, human or “paper” methods of contact tracing remain superior (Ebbers and Wijngaert 2020). What is more, app adoption does not appear related to time in market, suggesting that this is not merely a matter of time before adoption reaches critical thresholds.

Figure 5, then, should not be overly surprising. This figure plots the average level of openness in a country as a function of its app uptake. If contact tracing apps would effectively act as a substitute for restrictions on economic activity, then we would expect to see countries increasingly clustering in the top right hand and bottom left hand corners of the graph. There is some evidence of this for countries that are about 25% in app download rates, but across our entire sample, we see no clear effect. Governments are not linking economic openness policy to technology, likely because they cannot count on the latter to support the former.

In the next section, we explore in more depth why citizens are unwilling to adopt these apps.

Why don't citizens use contact tracing apps?

Our analysis in this section is interested in the Canadians COVID Alert app. This app is presently in use in 8 of 10 Canadian provinces. It is built on the Apple-Android infrastructure. It relies on Bluetooth signals to determine the proximity between two phones, which are then matched using anonymous identification codes on a users' phone, which are then stored centrally. This is regarded as the most privacy preserving app option on offer in the countries under examination (Ahmed et al 2020, Yasaka 2020).

Understanding individual-level decisions to download the COVID Alert app

Why do some individuals download the COVID Alert app? And what could persuade those who have not yet done it to do so? We begin by presenting an empirical model aimed at predicting app adoption. We then consider only those survey respondents who have not yet downloaded the app but are candidates for doing so (ie. they are smart phone users). We present these respondents with various experimental treatments within a survey and estimate the effect of these on their self-reported likelihood of adoption. By examining which treatments do and do not affect self-reported adoption likelihood, we can generate some insights into why a majority of the population is still yet to download and use the COVID Alert application.

Beginning with app adoption, we estimate a model that considers four sets of factors:

- Observable demographics, namely age, income, and gender.
- Beliefs about COVID, namely respondents' levels of concern about the virus and pandemic, as well as how long they expect the pandemic to last for.
- We include two sets of beliefs about digital privacy. The first is a well-established instrument which captures the degree of concern respondents have about the security of online data generally (CITE). The second is a new survey battery that captures the degree to which respondents believe the government should deploy technology to address the coronavirus pandemic.
- Finally, we include two items to measure government approval and political ideology.

As a benchmark, our data show that only one-third of respondents indicate that they have downloaded the app. What are the drivers then of uptake? Beginning with demographics, the model suggests that app adoption likelihood is declining in a respondent's age, and unrelated to their gender. There is no discernible relationship between a respondent's income and their likelihood of app adoption.

We also measure respondents' overall concern about the coronavirus pandemic, as well as their beliefs about how long preventative actions like social distancing, self-isolation, and

mandatory quarantines will need to continue. Our findings suggest that overall concern for the covid pandemic is strongly correlated with app uptake. Indeed, the likelihood of app adoption is 12 points higher among those who have the highest versus the lowest level of concern. Given that only one-third of respondents have downloaded the app, this is a substantively large effect. Contrarily, how long individuals expect other measures to be required is, however, not related to app adoption. This perhaps suggests that individuals do not see app adoption as a substitute for the collective need for other prophylactic measures, at least in the short term.

The next set of variables consider beliefs about technology. The results suggest that a generalized concern over data safety is not correlated with app uptake, at least not with any statistical precision. However, we find a clear and substantively large effect according to the belief that the government should use available technologies to combat the coronavirus. Those who hold this belief at a level above the median belief are 63 points more likely to have taken up the COVID Alert app.

It is reasonable to note here that there is likely an endogeneity between app adoption and belief that the government should use technology, such that those who have - for whatever welter of reasons - downloaded the app come to believe that the government should be relying on such tools. It is not possible for me to adjudicate causal ordering here, but we do wish to note this, because this belief will prove useful later on for figuring out how more uptake can be encouraged.

Finally, we find no effect for political variables. Overall evaluations of the federal governments' handling of the COVID pandemic are unrelated to app uptake. What is more, we find no differences among partisans in their likelihood to have taken up the app. This is not necessarily surprising, particularly given the broad, cross-partisan consensus on the need to fight the COVID pandemic (Merkley et al 2020). But it does foreclose the objection that app adoption is low because of polarized opinion on the governments managing the pandemic. Instead, it is low despite there being widespread agreement with and approval of governments' approaches to the virus.

Taken together, then, we have a reasonable model of adoption, in which individuals' choices to adopt the app are a function of both specific concerns about COVID, but also slightly more generalized beliefs about government usage of technology.

These findings complement previous work on contact tracing adoption. In an earlier analysis Loewen conducted in May 2020 in partnership with several Canadian Senators, we found widespread in principle support for a contact tracing app, in particular in support of a reporting function for public health officials. Concerns over government usage of data were a notable constraint on support for the app. At that point in time, however, an app had been neither

released nor promoted, so shifts in opinion since that time - not to mention substantial shifts in experiences of COVID - could have led to substantial changes in the correlates of app uptake.

Rheault and Musulan (2020) conducted a similar study of Canadians. They forward three key findings. First, by making individuals think about others avoiding social distancing, they find an increased support for contact tracing apps, suggesting that individuals have internalized the trade off between technology adoption and other prophylactic measures, or have at least come to understand contact tracing apps as a sort of safety measure. Second, they similarly find that anxiety about COVID predicts support for the app. Third, while they find substantial support for the app, they nonetheless find that resistance among the public is rooted in substantial concerns about violations of rights and freedom.

Horvath et al (2020) considered the preferences of citizens in the United Kingdom for various configurations of contact tracing apps. Providing survey respondents with choices between two contact tracing apps with randomly varied features, they find that citizens actually prefer a centralized system of data collection - in the hands of the trusted National Health Service - than a more decentralized system. Second, when they compare human contact tracing to digital contact tracing, they do not find clear differences, even when the potential of privacy breaches is emphasized for one method or technology over another. The takeaway here seems to be that trust in the management of data - especially as provided by a large, venerated, central provider - is key to supporting new technology deployment by a government.

Having some handle on which features lead people to take up a contact tracing app, we also wanted to explore more deeply what primes or interventions might convince those who have not adopted the app to indicate a greater willingness to do so. Using the same data from the Media Ecosystem Observatory, we exposed subjects to two experiments. Our experiments were limited to those respondents who have a smart phone but had not yet downloaded an app. In the first experiment, we provided them with different possible effects of adopting the COVID Alert app. Namely, we randomized the types of people who could be helped by individuals downloading and using the app. In the second experiment, we randomized subjects to receive information on how many people were using the app already, to see if greater population level engagement would itself spur greater adoption. we discuss the results of each experiment in turn.

Experiment 1: Effects on app adoption of helping different populations

In our first experiment, we randomized respondents to different information about who would be helped by a respondent downloading the COVID Alert app. In particular, in a baseline condition, respondents were told:

(Baseline) "Experts agree that if enough people download the app, it can help stop the spread of the virus. Knowing this, how likely are you to download and use the app in the next two weeks?"

Three other conditions randomly vary those who will be helped by downloads of the app. They were as follows:

(Underlying conditions) Experts agree that if enough people download the app, it can help stop the spread of the virus. This will especially help those with underlying conditions. Knowing this, how likely are you to download and use the app in the next two weeks?

(Older Canadians) Experts agree that if enough people download the app, it can help stop the spread of the virus. This will especially help older Canadians. Knowing this, how likely are you to download and use the app in the next two weeks?

(Help self) Experts agree that if enough people download the app, it can help stop the spread of the virus. This will especially help those who use the app. Knowing this, how likely are you to download and use the app in the next two weeks?

Subjects in every condition were given four response options: Very likely, somewhat likely, not very likely, and not at all likely. we have reshaped this to run from 0 to 1, where 0 indicates the lowest likelihood and 1 indicates the highest.

The goal here is to understand if more prosocial appeals, especially those that emphasize the vulnerable and the elderly can increase likely uptake, or alternatively whether emphasizing the benefit to the user instead motivates more uptake. We adjudicate between these possibilities by drawing group level comparisons to the control or baseline condition, which contained no information about who would benefit most from app uptake.

Our results suggest that the overall rate of willingness to adopt is relatively low, sitting somewhere between the survey categories of not very likely and somewhat likely. Most importantly, there is no appreciable difference according to variation in the target group who will apparently benefit most from the adoption of an app. we confirm these results in a regression, Table 1, Model 3, at the end of the paper, which includes the control variables we previously used to model uptake.

We further examined whether there is a conditional relationship between treatment and belief that government should use technologies to combat the coronavirus. The logic here is that citizens may have to first have a baseline belief that government should use technology to manage the coronavirus epidemic before they are willing to consider various appeals about

who will be most helped by COVID Alert app. The results are important. First, there is clearly a difference in willingness to adopt the app conditional upon believing government should use technology in the pandemic. Among those who hold this view, the average score across all treatments is 0.53. Among those who do not hold this view, it is just 0.29, a statistically significant difference (t=-11.85, p=.00). The treatments exercise no differential effect among those who believe government should use technology. Among those who do not hold this belief, however, the treatments do have a small effect, such that those who receive a treatment about helping older Canadians have significantly lower adoption rates than those in the baseline condition (t=2.26, p=.00). Nonetheless, the bottom line here is that there is little difference in potential uptake according to capacity of app adoption to help others. The action instead depends on whether citizens think the government should be using technology to fight the pandemic. These results are confirmed in Table 1, Model 4.

Experiment 2: Effects on app adoption of varying current uptake rates

For our second experiment, we wished to explore whether app uptake is conditional upon other individuals using an app. There are at least two explanations for why users would use an app conditional only on some other number of people adopting the app. First, supposing there is some cost to using the app (for example, through the potential invasion of privacy or via the costs of a false positive alert, which then sends the user into quarantine until an often-delayed test is returned), individuals may require some minimum number of other users to be using the app, such that the benefits of collective adoption outweigh the individual-level costs (note that these benefits could be conceived of as individual or pro-social). Second, individuals may determine the efficacy or otherwise attractiveness of the app based on how many others use it. Such an inference of quality according to usage rates is a common human practice.

To explore these possibilities, we ran a second experiment in which individuals were told about different (hypothetical) levels of app usage in their province. The exact wording was:

Suppose X percent of people in your province downloaded the app.
Knowing this, how likely are you to download and use the app in the next two weeks?

X was randomized with equal likelihood to take one of the following values: 1, 10, 20, 30, 40, 50, 60, 70, 80, 90.

Table 1, Model 4 estimates the (continuous) effect of this treatment, controlling for the other factors we have identified above. The effect is clearly statistically significant, but substantively underwhelming. Across the full range of the experimental treatment (1% population uptake to 90% population uptake), the estimated uptake likelihood increases from .07 to 0.16. There is an important heterogeneity, however, which is revealed when we condition the effect of the

treatment on the belief that the government should use technology to manage the COVID-19 pandemic. Table 1, Model 5 estimates the effect of the treatment conditional on this belief. For ease of interpretation, Figure 6 presents the estimated uptake rate across treatments for those who do and do not hold this belief. Among those who do not believe the government should deploy technology to confront the pandemic, hypothetical population level user rates have no effect on their likelihood to adopt the app. Contrarily, among those who do hold this belief, treatment rates are substantial. Indeed, the estimated likelihood increases from 0.41 in the 1% treatment condition to 0.60 in the 90% treatment condition.

There is an important lesson here. In addition to social distancing and other preventative measures against the coronavirus, citizens are willing to make use of government technologies, but this appears to be conditional on two factors. First, they have to believe that the government should deploy the technology. Second, they have to believe that others will use it. This is, in essence, a coordination problem conditional upon a collectively shared belief.

The challenge for governments is twofold, then. First, to convince citizens that technology can act as some partial substitute for other measures, and to convince them that the government should be able to use technologies which, on objective grounds, are less invasive and risky than those deployed continuously by technology companies. Second, they need to convince some threshold of citizens to take up these apps, with the aspiration that this will incentivize others to do the same. In the concluding section, we forward three related hypotheses for why governments have been, for the most part, unable to do these two things effectively.

Why can governments regulate but not promote widespread technology adoption

The last thirty years have been witness to a technological revolution (or series of revolutions), with little equal in the preceding centuries. The combination of rapidly-increasing computing power, the universal spread of the internet, and the adoption of smartphones has meant that individuals are connected to a complete and dense network, with massive amounts of data flowing in and out from them.

The impact of the internet and related technologies has not been universal, however. While some parts of daily lives have been fundamentally altered and improved - think here of ease of travel and navigation, shopping, and knowledge search and acquisition - other sectors have been much slower. Perhaps no sector - on a large-scale - has been less effective at harnessing the internet than governments. Arguably, the lack of widespread, app-based monitoring of infection, guiding of behaviours, and control of movement by infected individuals is a principal example. Just when governments could most benefit from citizens employing an app, in few places in the world can a significant percentage of citizens be convinced to do so.

In this essay, we have looked to Canada as a case study to try to understand why governments have been unable to push large scale and effective adoption of contact tracing apps. And this despite both widespread (and very costly) regulatory actions by governments and impressive social compliance and sacrifice on the part of citizens. By my lights, there are three broad explanations or hypotheses for why widespread adoption has not occurred. They are not mutually exclusive.

First, there is the possibility that citizens do not trust the government to properly handle their data and to not violate their privacy. Absent trust that their data is secure, citizens will opt out of using an app. Second, citizens may have a general lack of confidence that any technology that is promoted by a government could be effective. This belief in general incompetence may lead citizens to opt out of using an app and perhaps take up other options to stave off exposure. Third, it may be that while citizens do trust government to use technology, and do not necessarily believe that government is incompetent, they nonetheless expect that others will not use the technology in sufficiently high numbers, and so they themselves opt out of participating in this collective action.

We think there are merits to each possibility, and on balance evidence supports each of these hypotheses. Beginning with the first hypothesis, in our wave of data collection conducted the previous week to our experiments, we included a shorter module on contact tracing apps, and asked respondents for their agreement or disagreement with the statement “I don’t trust the government to protect my personal information enough to use a tracing app.” Nearly half of all respondents (49%) agreed with this sentiment, and the data are clear that these views are more commonly held among those respondents who have not downloaded the app. This

comports as well with findings from other countries (Altmann et al 2020, Abeler et al 2020). We likewise find support for the second hypothesis. For example, when respondents are asked "If using an app would end school and business closures sooner, I would be willing to do so", more than half agree 55%. Taken together, these data suggest both that citizens do not feel that their data are secure, and that they are unlikely to do much good anyways. This is hardly a ringing endorsement of the idea of digital government. The third hypothesis is likewise supported by some of our additional data. In the wave used to generate the results above, we also asked individuals what percentage of Canadians they believed had already downloaded the app. The mean response was 27%, with those who had downloaded the app showing a slightly higher estimate (30%) than those who had not downloaded (25%). Importantly, these are not wildly inaccurate estimates, in that Canadians correctly believe that not enough people have actually downloaded the app to make it a very effective tool. The story here then is not necessarily one about the government not being able to create an app that works, but the government not being able to create a digital solution which will be widely-embraced. In the face of low rates of cooperation, citizens perform here as they have performed in multiple other cases of collective actions: they choose not to contribute and instead invite a heavier hand from the government (Merkley and Loewen 2020).

How far do these results extend to other domains in which governments may want to employ technology but have been reluctant? On the one hand, it is difficult to say as in a Canadian context there are not ready examples at hand of times when governments have tried to push large-scale adoption of an app for a high-profile problem gripping an entire population. That is to say, this pandemic is unique. In other cases, the government has moved to put services online - for example, tax filing or license renewals - and has done so with some success. That said, such examples have little application here, as they deal with services and those for which the government has a monopoly. If obliged to file taxes, why not do it online?

But there are other domains where the government has resisted or even failed to imagine applications to address large-scale public problems. For example, governments have a vested interest in populations which eat healthier and are less obese - they spend substantial funds on public health campaigns aimed at these ends. And certainly, there is no lack of applications aimed at helping individuals control their food consumption. Why isn't the government in this game? Or, at least, why isn't the government in the game of encouraging citizens to even just use an app created by a third party?

There is a missed opportunity here, particularly if we think about app adoption as a collective action problem, where citizens are unwilling to use an app unless they are convinced that a large number of others are using it. The solution is not one of making the app mandatory (though there is support for that in our samples). Instead, the government could increase the attractiveness of downloading the app by paying citizens to use it. By subsidizing cell phone bills - which could be done through direct cash payments, through payments to telcos, or by

forcing telcos via regulation to cut the bills of those who download and use an app and then sidepaying the telcos - the government could increase usage at the margin. This is not a small change, as it would involve the government, for one of the first times ever, paying citizens directly in exchange for effectively being monitored (or at least alerted). It would not be without problems, as it may pay some of those who would comply anyways, create a moral hazard for others, and require an effective usage monitoring system. But, on the other hand, it may be vastly less expensive than regulating and limiting whole industries and sectors to the degree that many verge on shutting down.

The COVID pandemic has demonstrated that governments - at least in Canada - are ready to regulate again. They are not ready, however, to harness the power offered by technology to monitor and control the spread of the coronavirus. At a time when government matters more than any other time in decades, it has picked up worn out tools, and left newer ones on the shelf. This is in good part because citizens do not expect much else of their governments. This tells us a lot, we think, about the potential for the government to use technology into the future.

	Model 1 - uptake among all Canadians	Model 2 - adoption as a function of target groups	Model 3 - adoption as a function of target groups	Model 4 - adoption as a function of total usage rate	Model 5 - adoption as a function of total usage rate
Age	-0.18 (.01)	-0.04 (.59)	-0.04 (.52)	0.01 (.86)	0.02 (.78)
Gender	0.01 (.65)	-0.05 (.03)	-0.06 (.02)	-0.05 (.05)	-0.05 (.05)
Income	0.22 (.00)	0.12 (0.03)	0.12 (.04)	0.04 (.47)	0.04 (.47)
Covid concern	0.15 (.00)	0.31 (.00)	0.31 (.00)	0.26 (.00)	0.09 (.14)
Pandemic length estimate	0.06 (.29)	-0.05 (.31)	-0.06 (.17)	-0.03 (.54)	-.03 (.57)
Concerns for digital privacy	-0.12 (.07)	0.07 (.24)	0.07 (.22)	0.09 (.13)	0.09 (.14)
Belief in government using data	0.16 (.00)	0.17 (.00)	0.09 (.06)	0.18 (.00)	0.11 (.02)
Government approval	0.02 (.66)	0.05 (.29)	0.06 (.21)	0.05 (.30)	0.04 (.31)
Treatment 1a: help those with underlying conditions		0.04 (.29)	-0.04 (.45)		
Treatment 1b: help older Canadians		0.01 (.74)	-0.07 (.18)		
Treatment 1c: help app users		0.03 (0.44)	0.00 (.93)		
Treatment 1a* Government using data			0.14 (.06)		
Treatment 1b* Government using data			0.14 (.04)		
Treatment 1c* Government using data			0.03 (.57)		
Treatment 2: percentage of individuals using app				0.12 (.00)	0.04 (.55)
Treatment 2 * Government using data					0.15 (.09)
Liberal PID	0.06 (.18)	0.13 (.00)	0.13 (.00)	0.09 (.01)	0.09 (.01)
Conservative PID	-0.00 (.98)	0.01 (.70)	0.01 (.73)	0.01 (.86)	0.00 (.93)
NDP PID	0.07 (.23)	0,02 (.71)	0,02 (.69)	-0.02 (.72)	-0.02 (.72)
BQ PID	0.11 (.08)	0.05 (.37)	0.05 (.43)	0.02 (.75)	0.02 (.74)
Green ID	-0.06 (.23)	0.05 (.49)	0.05 (.43)	0.03 (.70)	0.03 (.71)
Other ID	-0.07 (.70)	-0.35 (.14)	-0.23 (.17)	-0.18 (.30)	-0.17 (.30)
Intercept	0.05 (.66)	-0.03 (.82)	0.03 (.78)	0.09 (.41)	0.12 (.31)

Table 1. Models of uptake and adoption. Every model is an OLS regression. The Ns are 1060, 543, 543, 543, and 543, respectively. Adjusted r2 measures are .06, .25, .26, .22, and .22. Each model includes fixed effects for province. Respondents from Alberta and BC are excluded, as the COVID Alert app is not presently used in those jurisdictions.

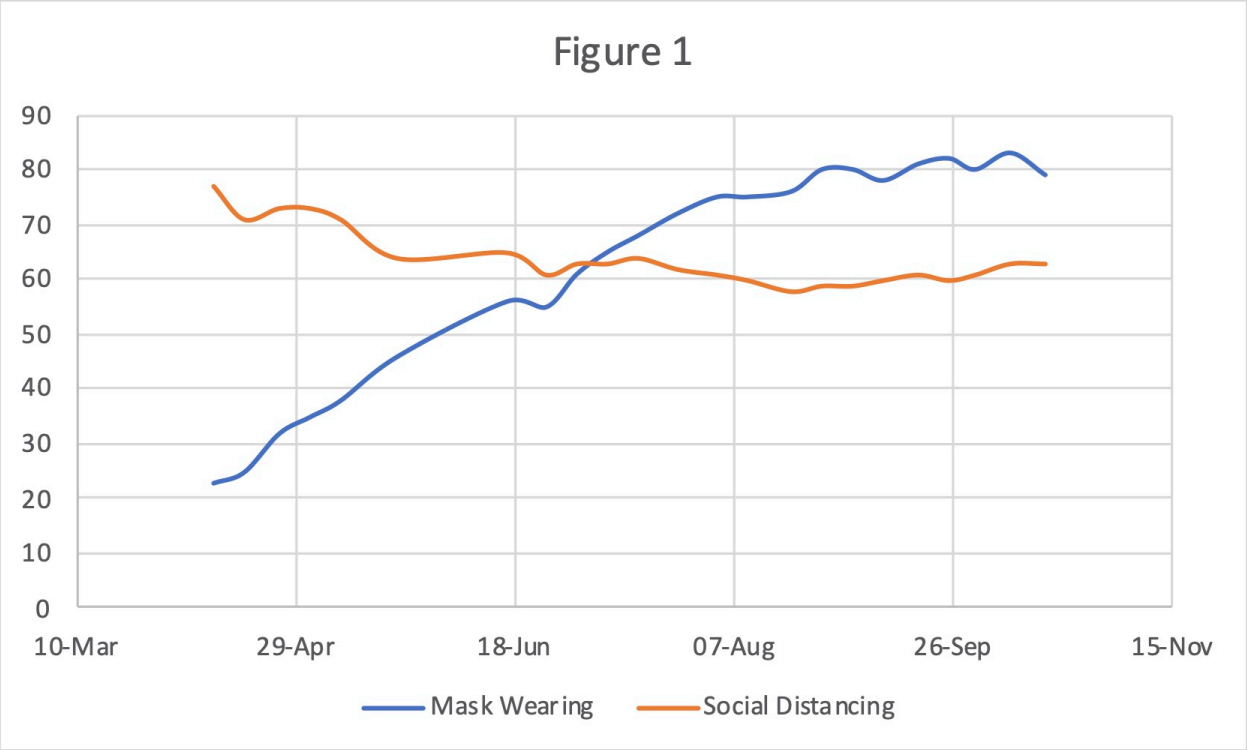

Figure 1: Social distancing and masking compliance in Canada (MEO data)

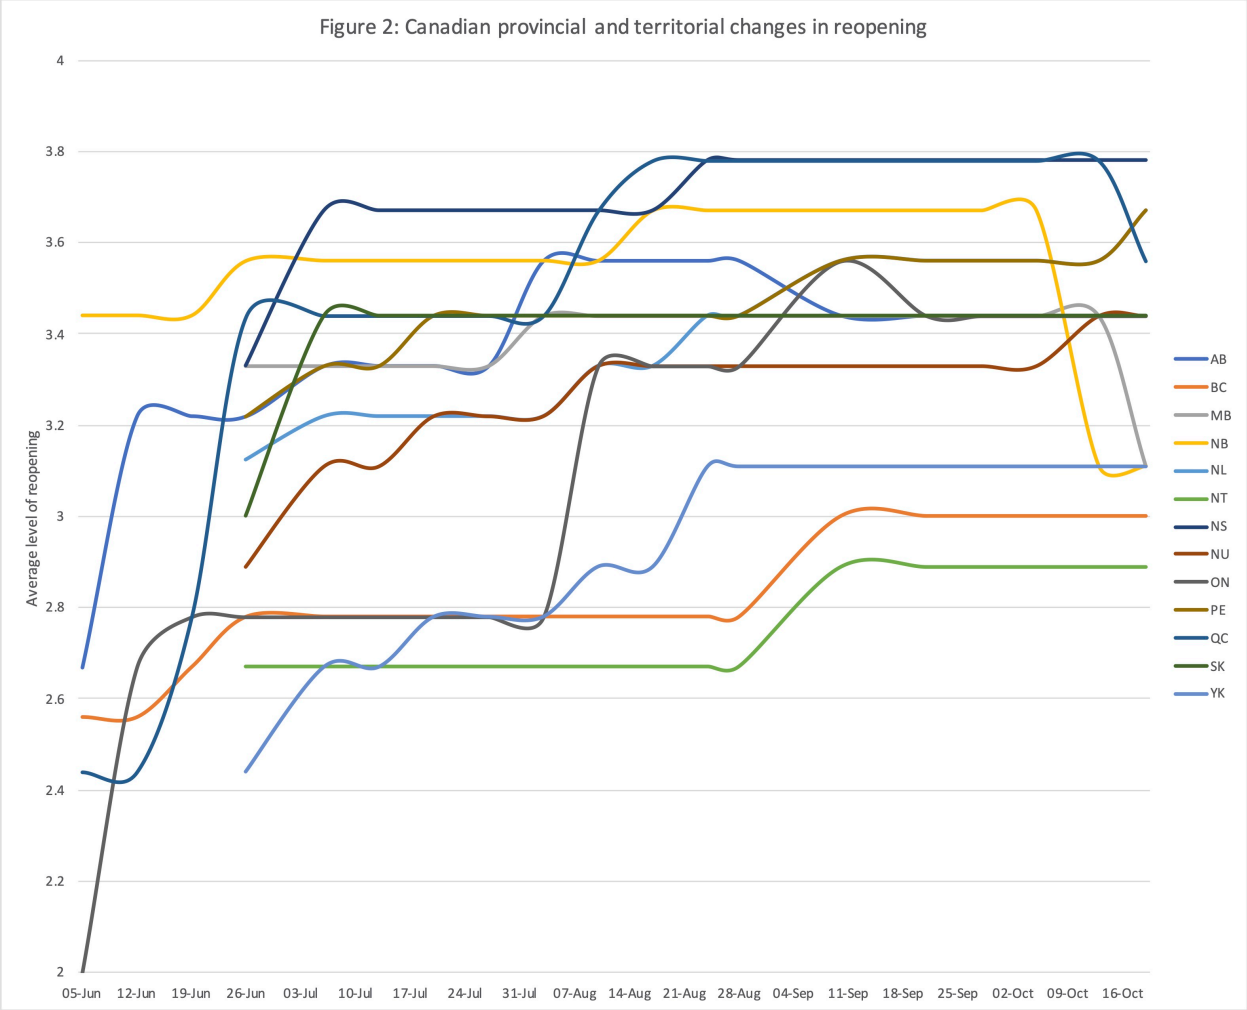

Figure 2: Canadian provincial and territorial changes in reopening (PEARL reopening data)

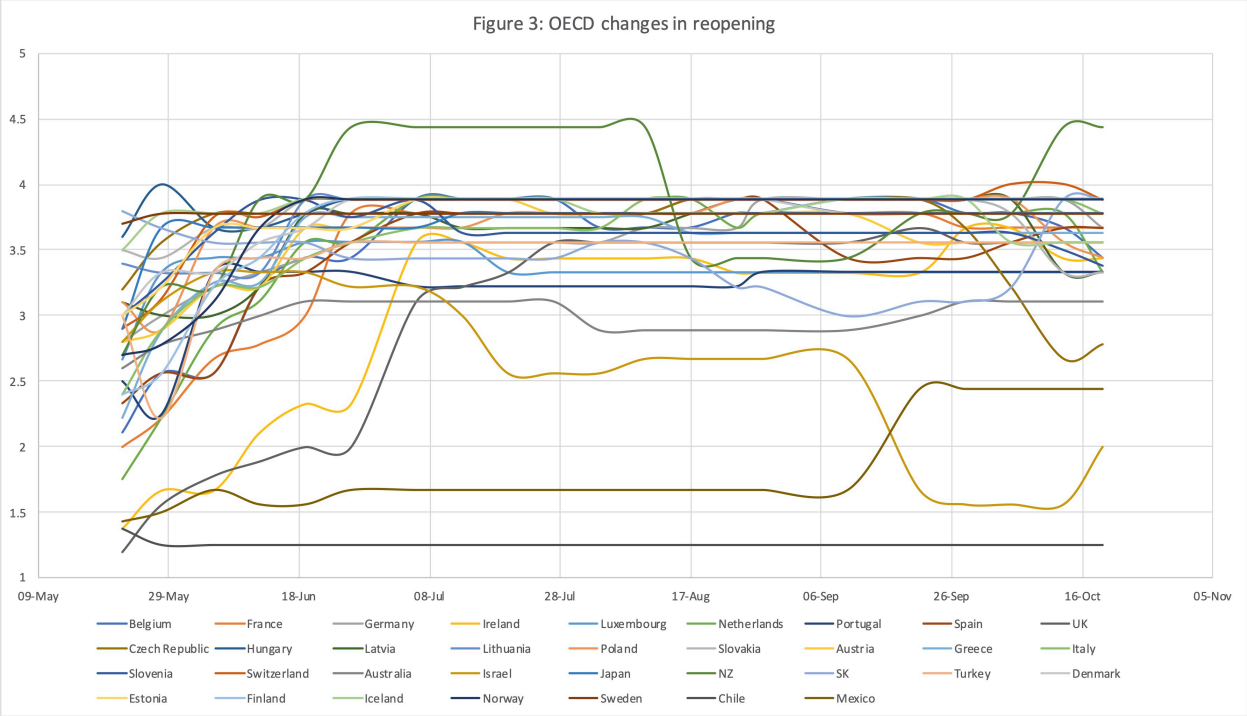

Figure 3: OECD changes in reopening (PEARL reopening data)

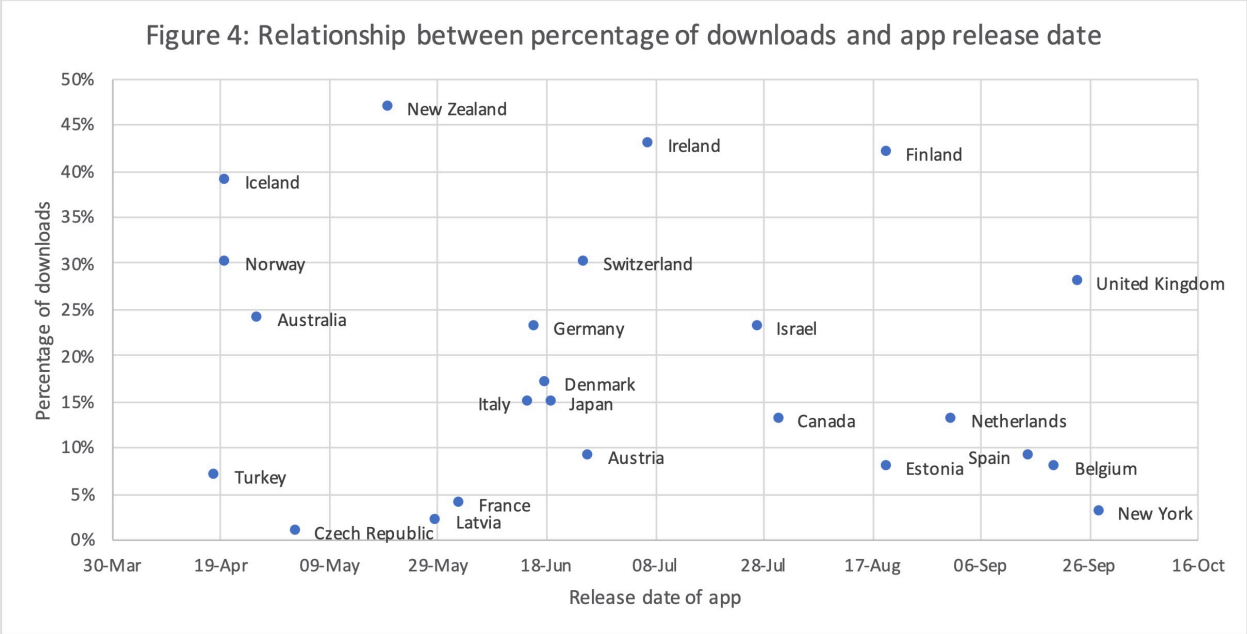

Figure 4: Relationship between percentage of downloads and app release date

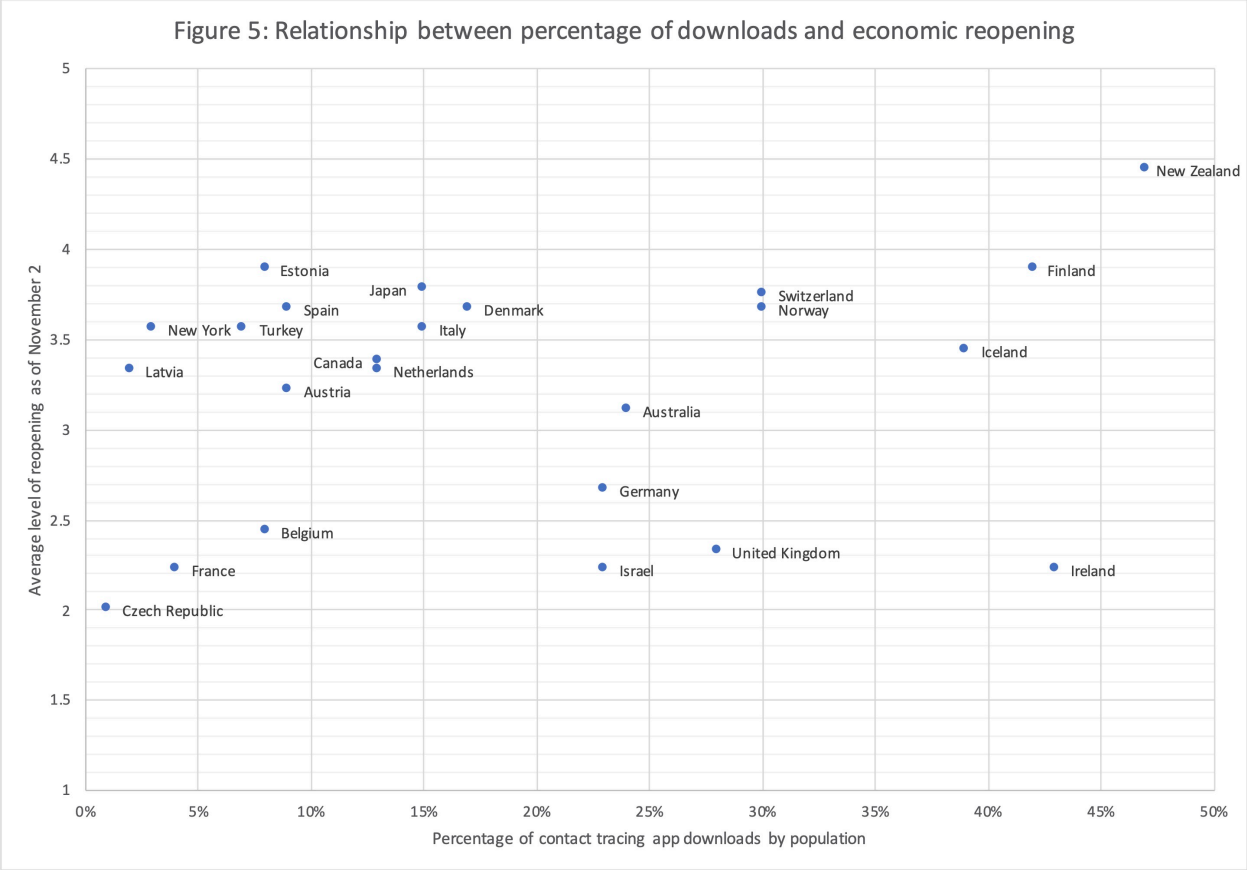

Figure 5: Relationship between percentage of downloads and economic reopening

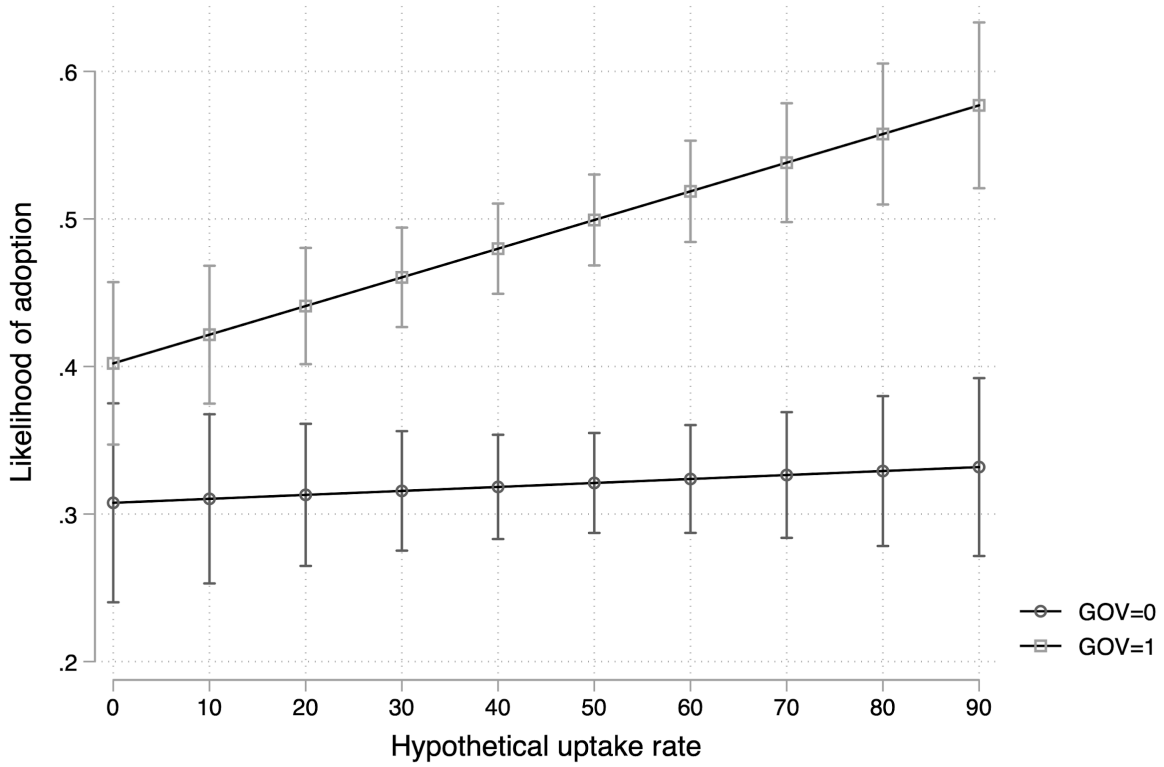

Figure 6: likelihood of app adoption according to hypothetical uptake rates by other citizens. Estimates are derived from Model 5 in Table 1 using the STATA margins command.

References

Abeler, J., Bäcker, M., Buermeyer, U. and Zillessen, H., 2020. COVID-19 contact tracing and data protection can go together. JMIR mHealth and uHealth, 8(4), p.e19359.

Ahmed, N., Michelin, R.A., Xue, W., Ruj, S., Malaney, R., Kanhere, S.S., Seneviratne, A., Hu, W., Janicke, H. and Jha, S.K., 2020. A survey of covid-19 contact tracing apps. IEEE Access, 8, pp.134577-134601.

Altmann, S., Milsom, L., Zillessen, H., Blasone, R., Gerdon, F., Bach, R., Kreuter, F., Nosenzo, D., Toussaert, S. and Abeler, J., 2020. Acceptability of app-based contact tracing for COVID-19: Cross-country survey evidence. Available at SSRN 3590505.

Bol, D., Giani, M., Blais, A. and Loewen, P.J., 2020. The effect of COVID-19 lockdowns on political support: Some good news for democracy?. European Journal of Political Research.

Christakis, N. Apollo's Arrow: The Profound and Enduring Impact of Coronavirus on the Way We Live. Little, Brown.

Ebbers, W.E. and van de Wijngaert, L.A., 2020. Paper beats ping: On the effect of an increasing separation of notification and content due to digitization of government communication. Government Information Quarterly, 37(1), p.101396.

Ferretti, L., Wymant, C., Kendall, M., Zhao, L., Nurtay, A., Abeler-Dörner, L., Parker, M., Bonsall, D. and Fraser, C., 2020. Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. Science, 368(6491).

Hinch, R., Probert, W., Nurtay, A., Kendall, M., Wymant, C., Hall, M. and Fraser, C., 2020. Effective configurations of a digital contact tracing app: A report to NHSX. en. In:(Apr. 2020). Available here. url: https://github. com/BDI-pathogens/covid-19_instant_tracing/blob/master/Report.

Horvath, L., Banducci, S. and James, O., 2020. Citizens' Attitudes to Contact Tracing Apps. Journal of Experimental Political Science, pp.1-27.

Johnson, B., 2020. Nearly 40% of Icelanders are using a Covid app—and it hasn't helped much.

Merkley, E., Bridgman, A., Loewen, P.J., Owen, T., Ruths, D. and 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/Revue canadienne de science politique, pp.1-8.

Merkley, E. and Loewen, P.J., 2020. Prospective Economic Costs Undermine Expectations of Social Distancing.

Rheault, L. and Musulan, A., 2020. Explaining Support for COVID-19 Cell Phone Contact Tracing.

Teichmann, L., Nossek, S., Bridgman, A., Loewen, P., Owen, T., Ruths, D. and Zhilin, O., 2020. Public Health Communication and Engagement on Social Media during the COVID-19 Pandemic.

Yasaka, T.M., Lehrich, B.M. and Sahyouni, R., 2020. Peer-to-Peer contact tracing: development of a privacy-preserving smartphone app. JMIR mHealth and uHealth, 8(4), p.e18936.

Appendix A - Variables used in regression analysis

Uptake - this is measured with the following question "Have you downloaded the COVID-19 Alert application for your smartphone?". The question was only asked of those who indicate that they own a smartphone. Non-smartphone users and non-adopters are coded 0, while adopters are coded 1.

Adoption - this is the likelihood of app adoption by those who have not yet adopted the app. The question wording was "How likely are you to download and use the app in the next two weeks?" Response categories were Very likely, somewhat likely, not very likely, and not at all likely. The variable is reshaped 0 to 1, with 1 indicating the highest likelihood.

Age is measured in years, rescaled from 0 to 1.

Female reads 0 if a respondent identifies as male and 1 if a respondent identifies as female.

COVID concern - this the average of three statements: "How concerned are you about the coronavirus pandemic?" (Very concerned, Somewhat concerned, A little concerned, not at all concerned); "How serious of a threat do you think the coronavirus (COVID-19) is to yourself?" (Very serious, somewhat serious, Not very serious, Not serious at all); "How serious of a threat do you think the coronavirus (COVID-19) is to Canadians? (Very serious, somewhat serious, Not very serious, Not serious at all). Variable is reshaped from 0-1, with 1 indicating greatest average concern.

Pandemic length - this variable captures who long individuals believe various measures will need to remain in place. The question was: "For how long do you think it will be necessary to keep each measure in place to control COVID-19?". It includes each of the following measures : "Social distancing and good hand hygiene", "Self-isolation — voluntary stay at home orders for all those experiencing symptoms." Quarantine — Mandatory stay at home orders for all those who have tested positive." Response categories were "End immediately, less than 1 month, 1-6 months, 6-12 months, 12-24 months, until a vaccine is found." The average response across all three actions was taken, and reshaped from 0 to 1, with 1 indicating the most amount of time.

Government support - Government support is the average of three measures of support for governments' handling of the COVID pandemic. The question wording is: "To what extent do you approve or disapprove of the (federal/provincial/local) government's handling of the coronavirus pandemic so far?" Response categories were strongly approve, somewhat approve, neither approve nor disapprove, somewhat disapprove, and strongly disapprove. The variable is reshaped from 0 to 1, with 1 indicating the highest support.

Concerns for digital privacy - This is a measure taken as the average of the following five questions. "In general, how concerned about you about the following": "Your privacy while you are using the internet"; "Online organizations not being who they claim they are"; "Online identity theft."; "People online not being who they say they are"; "That your credit card may be intercepted by someone else when you use it to buy something on the internet". Response categories are not at all, not very much, somewhat concerned, very concerned. Higher scores indicate more concern. All questions were averaged and reshaped 0 to 1.

Government using data - This is a measure that takes the average response to three questions: "I am more concerned about government having my data than private companies" (reverse coded); "I am willing to give up some privacy in order to slow down the spread of the coronavirus"; and, "Government should use as much technology as necessary to slow the spread of coronavirus". Response categories were strongly agree, somewhat agree, neither agree nor disagree, somewhat disagree, strongly disagree. The variable is split at its median, with those having above median agreement coded 1 and others 0.

Partisan identifications were measured using the standard question of "In federal politics, do you usually think of yourself as a(n): Liberal, Conservative, NDP, Bloc Quebecois, Green, Another party, no party, Don't know."