Surveillance & Privacy - January 2021

# *COVID Alert's Privacy Promises and Surveillance Risks*

Ignacio N. Cofone


Centre for MEDIA, TECHNOLOGY and DEMOCRACY
McGill UNIVERSITY
MAX BELL SCHOOL of PUBLIC POLICY

# Contents

# About the Author

## IGNACIO N. COFONE

Assistant Professor and Norton Rose Fulbright Faculty Scholar, McGill University Faculty of Law

Ignacio Cofone is an Assistant Professor at the Faculty of Law at McGill University where he teaches courses in privacy law, business associations, and artificial Intelligence law. Dr. Cofone's research explores the ways in which law should adapt to technological and social change, focusing on privacy and algorithmic decision-making. Before joining the McGill faculty, he was a research fellow at the NYU Information Law Institute, a resident fellow at the Yale Law School Information Society Project, and a legal advisor for the City of Buenos Aires. He holds a law degree from Austral University, an LLM and JSD from Yale Law School, an MA from the University of Bologna, and a joint Ph.D. from Erasmus University and Hamburg University where he served as an Erasmus Mundus Fellow.

*I thank Malaya Powers and Jeremy Wiener for their fantastic research assistance and gratefully acknowledge research support from the Social Sciences and Humanities Research Council (SSHRC).*

ignacio.cofone@mcgill.ca.

# About the Series

In this essay series, *Watching the Watchers: The New Frontier of Privacy and Surveillance under COVID-19*, McGill's **Centre for Media, Technology and Democracy** explores the policy, legal and ethical issues of (new) surveillance tactics in times of crisis.

In the wake of the 2020 global pandemic, governments and corporations around the world are adopting unprecedented data-gathering practices to both stop the spread of COVID-19 and transition to safer and more economically stable futures. This essay series examines how public and private actors are using pandemic response technologies to capitalize on this extraordinary moment of upheaval. It convenes a diverse group of experts to examine the policy, legal, and ethical challenges posed by the use of tactics that surveil and control populations around the world. With a focus on wide-ranging topics such as cybersecurity, racial justice, and worker surveillance, among others, this series offers a roadmap as policymakers confront the privacy and human rights impacts of crises like the novel coronavirus in the years to come.

# Introduction



On 31 July 2020, the federal government rolled out a contact-tracing app called COVID Alert to help public health agencies across Canada contain the spread of COVID-19. The app was presented by the Prime Minister as a tool that can enhance manual contact tracing, which has limits in its effectiveness to track exposure when operating alone.[1]

Due to constitutional constraints related to federalism (provinces, not the federal government, have authority over healthcare in Canada), it's each province's decision whether to incorporate the app. The app was first rolled out in Ontario exclusively before doing so nationally.[2] Being nationally available, some provinces, like Quebec, agreed to its incorporation as recently as last month. Others, including the territories, BC, and Alberta, are still deciding whether to do so.

The app serves an important role in preventing the spread of the virus and containing the pandemic, enhancing the ability to contact-trace enormously compared to manual contact tracing. Contact tracing is a longstanding method to contain widespread contagious disease outbreaks. For diseases as contagious as COVID-19, contact tracing serves an important role in containment.[3] This is the first time that the world meets a pandemic with the ability to use surveillance technology to leverage these tracing efforts.

But many activists and organizations have warned about the dangers of COVID Alert and other contact-tracing apps, particularly regarding the risks to human rights that can accrue from their ensuing surveillance.[4] The apps enable governments and private companies across the world to track and surveil citizens, and often involve combining highly sensitive information such as health information and location. While the sensitivity of health information may be evident to most people, location is equally revealing. Location data reveals "highly sensitive data about people's behaviors, patterns, and personal

life."[5] It not only reveals where you are but also what establishments you go to, who you spend time with, when and for how long you do so, what kind of activities you engage in, among other information about you.[6]

COVID Alert handles many of these risks well. But it's impossible for it to solve all of them while remaining functional. Risks thus remain. This doesn't mean that people shouldn't use the app, but it does mean that the app has drawbacks and limitations that prevent it from being a holy grail for containing the spread of the pandemic. It also means that it's productive for both individuals and policymakers to have a clear picture of the resolved and remaining drawbacks when making individual and policy decisions. This short piece aims to provide an overview of the app and of the key privacy-related risks that, like other contact-tracing apps, it must address.

# Canada's COVID Alert app

COVID Alert was built on a new Application Programming Interface (API) from Apple and Google, which acts as a framework for developers to build apps that execute proximity tracing. COVID Alert executes this with open source code for exposure notification developed by Shopify, which collaborated with government organisms including Health Canada, Innovation, Science and Economic Development Canada, the Canadian Digital Service, and the Ontario Digital Service.[7] Each Canadian resident can choose whether to download the app. If they download it and later test positive for COVID-19, they can choose to anonymously input the positive result in the app through a dedicated one-time code (the "key"). In more recent updates, those who test positive can also enter the date in which they were tested and the time that their symptoms started.[8]

The app will then notify those who were recently in close contact with that person and have also downloaded the app.[9] Close contact is defined as spending more than 15 minutes within 2 meters and recently is defined as within 14 days. In one version of the app, it says "You've been exposed in the last 14 days. Someone you've been near has reported a COVID-19 diagnosis through the app. You were close to them for 15 minutes or more. You're at risk of being infected."[10]

The app works through Bluetooth. It taps on Bluetooth's ability to broadcast a unique identifier for each device called a MAC address. Because Bluetooth can project identifiers within a certain range (as you may have noticed if you walk far away from your phone with your Bluetooth headphones on), it can tell if two devices were in proximity (within the Bluetooth range) for any amount of time.[11] Instead of sharing MAC identifiers among devices, COVID Alert uses this technology to exchange "tokens," which are temporary identifiers generated by the device.[12] A device's token (identifier) changes every 15 minutes and is stored locally in the person's phone.[13] The device keeps two lists: one for the tokens it generated and one for the tokens it came into contact with recently.

When someone inputs a positive result into the app, they can allow the app to share their recent tokens with a central server and those who were within range.[14] A key advantage of this system is that it avoids the need for public health workers to access testing records and reach out to individuals through phone or email. Another key advantage is that it doesn't rely on individuals' potentially feeble memory in listing to the manual tracer every person that they have been in contact with during the last 14 days.[15]

The Bluetooth token system avoids using GPS location and, more broadly, real-time tracking altogether. Google and Apple thus presented their system from the start as one that has "privacy and security core to design".[16]

# A brief taxonomy of contact-tracing apps

Like with anything else, there are many ways to sensibly classify contact-tracing apps. In how they perform on privacy, two distinctions seem most relevant: one regarding the tracing method and one regarding the storage method.[17]

The first is whether the app collects and provides proximity data or location data.[18] That is, whether the app works based on Bluetooth or GPS. While GPS reveals each device's location, Bluetooth doesn't reveal location, but only the devices' proximity with each other.[19] Bluetooth thus makes it more difficult to infer personal information.[20] While location data can easily be combined with data about where stores are and with other people's location data, it's difficult to make these inferences based on Bluetooth data. In other words, Bluetooth doesn't reveal a person's location history, and makes it much more difficult to infer other personal information. It's thus substantively less intrusive.[21]

The second is whether the app has localized or centralized storage.[22] That is, whether the information on spread (collected through Bluetooth or GPS) is stored in each device and shared only among devices, or is transmitted to a central location such as health authorities. Some people call localized storage apps "exposure notification apps" and they reserve the term "contact-tracing apps" for centralized storage apps.[23] They do so because only with centralized storage can health authorities map social networks, analyze spread trends, and study population movements. One could think that both types trace contact, just that one shares it with a centralized third party and the other one doesn't. But that act of sharing with a centralized third party is enormously consequential because it allows for the aggregation of information. That entails numerous benefits from an epidemiology perspective and, at the same time, exponentially increases privacy risks, such as the risk of data misuse.[24] Localized storage m makes it much more difficult to identify, track, or study individuals based on app data.[25]

We can combine these in a two-by-two table:

| | LOCAL STORAGE | CENTRAL STORAGE |
|---|---|---|
| Bluetooth | Most privacy sensitive | Medium privacy sensitive |
| GPS | Medium privacy sensitive | Least privacy sensitive |

In this classification, COVID Alert fits in the top-left quadrant, in the most privacy-sensitive type of contact-tracing apps. COVID Alert is probably the least privacy-invasive functional option that the government could have adopted. That means that it avoids many of the risks of other alternatives (and, conversely, it has fewer functionalities than them). It also means that, because of its functionality, surveillance risks remain. Some of these risks may be inevitable for the app to work.

The placement of an app in this taxonomy will affect its level of privacy risks. These risks exist because of the nature of the surveillance and the extent to which the app is privacy sensitive. This includes obvious risks such as re-identification (anonymity) and lack of consent, as well as less obvious ones such as inaccuracy risks and discrimination risks that exist because surveillance is often unevenly distributed, disproportionately affecting the most vulnerable. The following two sections detail what this means about how COVID Alert handles some of the privacy risks that contact-tracing apps can create.

> That act of sharing with a centralized third party is enormously consequential because it allows for the aggregation of information.
>
> That entails numerous benefits from an epidemiology perspective and, at the same time, exponentially increases privacy risks, such as the risk of data misuse.

# Risks in guarantees against surveillance

## a. THE TRICKS OF REDUCING SURVEILLANCE THROUGH CONSENT

To understand COVID Alert's measures to reduce surveillance—consent and anonymity—it's useful to understand when surveillance turns into a problem: the problem of privacy harms.

Revealing information isn't necessarily a bad thing. Revealing information about contagion that is shared with others who may have been exposed is exactly what we would hope contact tracing does. Privacy harm derived from surveillance is about revealing, together with this information, other information that could be harmful to the person, such as habits, preferences, or company.[26] Privacy harm is largely about disclosing, together with the relevant information, lots of other irrelevant information.[27] That is the surveillance that well-designed systems attempt to minimize.

This is relevant to contact-tracing apps. Location reveals plenty of other information that is irrelevant for contact tracing but can produce privacy harms. People's location data reveals, for examples, their movements, habits, and preferences. While Bluetooth doesn't provide location data (proximity data doesn't reveal location directly), it still produces a trail of people who the individual was in contact with, when, and for how long,[28] which may contain sensitive information. Even without location, proximity can reveal a lot of information: who you were with, when you were with them, and for how long. It's effectively a detailed map of anyone's social interactions. Depending on how many people you were in contact with, and on how many those people were in contact with, these data can be used to infer location data through what is called a linkage attack.[29] This risk is lower for COVID Alert compared to other contact-tracing apps because of the frequency with which the tokens change.

The first way in which surveillance is addressed in contact-tracing apps is consent. The use of COVID Alert is meant to be completely voluntary.  The underlying idea

is that, if a user reaches the conclusion that the apps' benefits don't outweigh the privacy risks that it represents for them, they won't download it.

The consent guarantee comes with a caveat and a limit. The caveat is that, in privacy, consent doesn't guarantee that people's rights won't be violated.[30] The limit is that, for any contact-tracing app, the extent to which consent expresses actual voluntariness is questionable.

Explaining how exactly the collection, use, and disclosure of these data work in a way that users will understand is genuinely difficult. For consent to be valid, each individual has to understand what each of these things means. Given the app's complexity and the number of unknowns about how the data could be aggregated and used, that is unlikely to happen. In fact, COVID Alert has minimal explanations for its users, other than explaining random identifiers and guaranteeing a concern for privacy.

Even if formal consent is achieved in an app that remains optional, it is difficult to ensure that the app isn't informally, effectively, mandatory. The app can turn effectively mandatory if, for example, stores make having installed the app a requirement to enter or other private parties condition services on it.  More problematically, an employer could make the app mandatory for its employees and use it to monitor if employees have been infected. For the app to be voluntary, the government should include guarantees that downloading it or using it won't be a condition for social participation or inclusion.[31]

There is an additional problem. If we pass those two hurdles, we have to ensure that the app is rolled out in a way that isn't coercive for the people who agreed to it. This could happen, for example if police action is made dependent on the app, such as by coupling data on who is infected with enforcement of quarantine or with eventual criminal sanctions. Relatedly, if the government or private actors place restrictions based on reporting, people may under-report. Coercion could also happen indirectly from private discrimination, described in the last section of this essay.

An emphasis on consent has a downside in terms of the app's effectiveness. For a contact-tracing app to work effectively, it needs a substantial percentage of the population agreeing to use it. This is more difficult to achieve if its use is truly voluntary. While the initial effectiveness requirement publicised in

> Even without location, proximity can reveal a lot of information: who you were with, when you were with them, and for how long.
>
> The caveat is that, in privacy, consent doesn't guarantee that people's rights won't be violated.

the media of "at least 60% of the population" was disproven as a misunderstanding of the original research,[32] low adoption makes the app less useful.[33] Particularly, low adoption leads to the problem of high inaccuracy described below (high false negatives) that may lead the app to be ineffective or even detrimental.[34] If only 20% of the population uses the app, then someone using the app having contact with a person who is infected and contagious has a lower chance of being alerted of such contact.[35] With 20% of the population adopting the app (a low estimate), even assuming perfect detection by the app, it would detect 4% of all contacts. In Canada, the 60% adoption flagged as the threshold for more effectiveness means that 22.554.000 people would have to download and use it.

### b. THE IMPOSSIBILITY OF ANONYMITY

To minimize unnecessary surveillance, most contact-tracing apps come with the promise to anonymize all personal information that they process. This initially included COVID Alert, which then removed anonymity promises after receiving feedback from the Office of the Privacy Commissioner (replacing it with more ambiguous language of it being privacy-safe).[36]

This change was an improvement. Promises of anonymity in these data are impossible to attain.[37] The issue is that perfect anonymization isn't possible because truly anonymous data doesn't really exist.[38] One can only remove personal identifiers (de-identify) data, like the app does when using tokens that change rather than people's name. But any data can be re-identified with enough effort.[39] Claiming that the data collected will be anonymous, and therefore not privacy-invasive, is thus a red herring.

The probability of re-identification of any de-identified data depends on the threat model. How likely data are to be re-identified depends on how much effort it needs and how valuable the prize of re-identifying it is. In the case of contact-

tracing apps, the data are easy to identify because it includes granular proximity or location data. And they are worth it for malicious actors to spend costs to obtain and reidentify it. In the case of location data paired with health data, the prize is sizeable. Because it combines sensitive health data with location, it's prime for third-party misuse. That makes COVID Alert's decentralized use of proximity data hugely important. But while proximity data isn't as easy to re-identify as location data, that doesn't make it fully anonymous. The oxymoronic use of the phrase "anonymous identifiers" is the most glaring piece of evidence of this fact.[40]

The re-identification risk is bigger for GPS than it is for Bluetooth and it's significantly bigger for centralized than for localized data storage. So COVID Alert is better than its alternatives on this. But the risk isn't inexistent with Bluetooth. Users can be identified if proximity data are combined with other data. This would be facilitated, for example, if public health agencies were to solicit postal codes upon creating an account for the app.  COVID Alert doesn't do this in its current version, and it states after being installed that "it has no way of knowing your location;" it only asks for your province, and it's optional. But even without explicitly requesting location, it's possible to collect one's IP address, thus knowing which Wi-Fi signals the phone joined, effectively making it into a location tracking app.[41] The agency running the server receiving the keys will inevitably be able to link keys to the IP that is uploading them and know who the individual corresponding to each key is.[42] This means that, despite anonymity guarantees, the government can always know who reported having tested positive for COVID-19 if it wanted to. This risk is inevitable.

These data are especially easy to re-identify with a lockdown and, to a lesser extent, with physical distancing. If someone who was in contact with three people during the 14-day period receives a notification, (absent false positives) she would have a fair estimation of who the infected person may be. It's easier to identify who triggered a positive notification the fewer people you see.

This makes it easier not only during lockdown but also easier for people who live in areas with a more dispersed population.[43]

Anonymizing these data isn't only impossible, but it would also not resolve surveillance even if possible. "Anonymous" data still reveals information

and trends about groups.[44] There is, thus, potential for data that remains de-identified to harm members of those groups by, for example, being used in a discriminatory manner, even unintentionally.[45] Harms in group privacy do not rely on re-identifying individuals: decisions can affect de-identified individuals on the basis of group attributes (such as gender, sexual orientation, political preference).[46] Breaches in group privacy can amount not only to discrimination, but also infringements of other Charter rights such as freedom of opinion and expression, and freedom of association.[47]

# Risks stemming from the distribution of surveillance

## a. TYPES OF INACCURACY AND TRUST

Algorithms work based on proxies. The effectiveness of any algorithm depends on the strength of the proxies it uses.[48] In this case, "having my device exchange tokens with a device that registered a positive test result" is a proxy for "having contracted COVID-19." Perhaps more importantly, Bluetooth signal strength is a proxy for shared airspace.[49] The timing of the tokens works to improve this proxy. "The 15 minutes of close contact guideline for possible transmission means that casually passing someone on the sidewalk isn't something an app would need to track and log."[50]

This is a decent proxy. In fact, it may be one of the most acceptable proxies available. But, like all proxies, it's imperfect, and it's important to pay attention to the distributional aspects of its imperfections. We should distinguish between two types of errors. False positives are errors that say you were exposed to COVID-19 when you were not. False negatives are errors that lead to mistakenly thinking that you haven't been exposed to COVID-19 when you have. Distinguishing between false positives and negatives is important because they're differently costly and they differently impact people's use of the app.

Bluetooth can tell how close you were to someone and for how long. But it doesn't know if you or them were wearing a mask. It also doesn't know if there was a plexiglass between you. It doesn't even know if you were each inside of your own car, with your windows closed, waiting at a red light next to each other. It doesn't know that there's a wall between your neighbour and you and you don't spend every night sleeping next to each other. It doesn't even know that the person who lives below or above you isn't in your living room. Obstacles like walls decrease the strength of Bluetooth signals, but they stop Bluetooth signals significantly less than they stop the virus' aerosol transmission. Face masks and hand sanitizer, as they reader may have guessed, reduce the likelihood of transmission but don't alter Bluetooth signal. These errors are what we call "false positives."

> To understand COVID Alert's measures to reduce surveillance—consent and anonymity—it's useful to understand when surveillance turns into a problem: the problem of privacy harms.

False negatives are less varied than false positives in COVID Alert. The more common false negatives would occur when being close to someone who doesn't use the app, doesn't get tested when presenting symptoms, or doesn't log a positive test result.

False positives may have an effect on people's mental health and businesses' finances.[51] "One or more intentionally false reports of positive tests could have financial and social consequences for businesses or people who either believe they are infected, or whom other people believe are infectious."[52] But more importantly, false positives matter because they erode trust in the app. If I get lots of false positives while being confident I'm negative, then I won't trust what the app tells me anymore, and the app's usefulness will be undermined. False negatives, in turn, may lead infected people to erroneously think that they're not infected, leading to them not getting timely treatment or failing to quarantine and spreading the virus to others. An app with many false negatives could be worse than no app for containment purposes if people ignore other signs of disease due to trust that they haven't been exposed based on app results.

The costliness of both of these errors depends on how the app is used. They become more problematic when an inadequate amount of trust is placed in the app. Cathy O'Neil explains this in terms of the importance of not placing blind faith in data outcomes, which she calls being a data skeptic.[53] The importance of this skepticism increases with high-stakes systems where the process is difficult to understand for users, such as contact-tracing apps. For these apps to be effective, the right amount of deference (not more, not less) must be given to the machine. And people often defer too much.

If someone doesn't trust the app's outcome, the whole system becomes unhelpful for them—algorithmic predictions are of little use if ignored. In contact tracing, "[t]he whole system depends on trust. If users don't trust that an app is working in their best interests, they will not use it."[54]

Placing too much reliance on the app is also problematic. Overreliance on the app when there is a false negative might make people less likely to take other health measures because they have received a negative from the system. Overreliance when there is a false positive can impose unnecessary stress on families and overburden the healthcare system. Similar problems in terms of false positives and false negatives arise if the ones over-relying are not private individuals but public health officials.

### b. MAGNIFYING INEQUALITY

Another problem with inaccuracy is that algorithmic error is rarely distributed evenly.[55] In contact-tracing apps, like in many other systems, inaccuracy is distributed in such a way that is likely to disproportionately harm the most vulnerable. Because of the importance of consent and of trust in the technology, this happens with groups that are over-surveiled and groups that are left out.

Contact tracing is most useful for the most vulnerable. Among them are the economically vulnerable. Construction workers, delivery people, domestic workers, cashiers, bus drivers, taxi or Uber drivers, constantly come in contact with customers or each other in doing their jobs. Those who don't have an office job that they can do from home, but rather have a job where they are in contact with lots of people, are exactly those that public health officers would want to trace and are exactly those who would want to be notified of exposure.

These groups need the app more, but more proximity to others leads to more likelihood of an app showing a positive result even when safety measures are taken to keep the real risk low. Those who don't have an office job or can't afford working from home are more likely to show a positive result on the app. They are thus more likely to bear the economic and psychological harms of false positives.

The app may also be less optional for them if requested by their employers or clients. At the same time, if positive results lead them to face negative outcomes at work (for example by having to take time off after receiving a positive alert),

then the app amplifies problems of economic inequality. Their situation pairs error distribution with the consent element I mentioned above. This fact feeds into a general narrative of marginalized communities having long been shown to disproportionately experience surveillance and its accompanying risk.[56]

Testing positive for COVID-19 may also carry a social stigma that is unlikely to dissipate. When receiving an alert, people can draw correct or incorrect inferences to place blame.[57] The ease of this exercise partly depends on the apps' design. If alerts upon possible exposure are immediate, it may be easy to guess who triggered it. If alerts accumulate and the person is informed of collected possible exposures after some amount of time, its ease depends on how long the delay is and people may be less confident in their estimations of blame. COVID Alert has some delay, sending notifications once a day, which is a desirable feature.

Publishing aggregated location data about COVID-19 outbreaks, which COVID

> In contact-tracing apps, like in many other systems, inaccuracy is distributed in such a way that is likely to disproportionately harm the most vulnerable.

Alert doesn't do, would have intensified this. Outbreaks being linked to certain communities, religious groups, or neighborhoods, would likely create or intensify stigma towards them, just like they did for HIV.[58] This is particularly important in a social context where COVID-19 has already intensified anti-Asian xenophobia and racism.

On the other end of the spectrum are the vulnerable that the app leaves out. Those who are more vulnerable to the virus because of their living conditions are less likely to own a smartphone (and, particularly, a smartphone new enough to support the app).[59] This isn't a coincidence. The elderly living

in care facilities, the homeless, migrants, refugees, and prisoners all live in conditions that make them more likely to contract the virus because they lack the means that will also give them access to a device that can support the app to detect or report it. This leads to an irony that has been pointed out for other technologies.[60] These particular vulnerable groups will receive less surveillance and thus less accuracy. This may seem like a good thing, but it's not. First, the lack of accuracy may mean that they disproportionately suffer from the side effects of uneven surveillance in terms of stigma, piling on existing inequalities. Second, to the extent that the public health response relies on surveillance, as opposed to, for example, a longer quarantine, they are left out of such effort.

In sum, health data that cannot be fully anonymous, cannot rely on consent to prevent harm, and is prone to stigma, combines in problematic ways with asymmetric rates of error for marginalized groups.

# Conclusion

This piece explored four interrelated risks that contact-tracing apps inevitably deal with due to the type of surveillance that they exert to be functional, together with how COVID Alert measures up in each of them. COVID Alert has robust security measures that make it better than most other alternatives. Risks, however, inevitably remain. These are risks that, conditional on unfolding the app, Canadian policymakers can and should be attentive to and try to curb, both at a federal and at a provincial level.

But these are also risks that users can keep in mind when deciding how to use the app. They can do so, for example, by selectively turning Bluetooth on and off to avoid disclosing information that they would rather not, by making sure they understand the working and consequences of the app before weighing the costs and benefits of using it, by placing the right level of trust in the app's output, and by being judicious about drawing inferences about others.

## Endnotes

1 Office of the Prime Minister. (June 18, 2020). *Prime Minister announces new mobile app to help notify Canadians of COVID-19 exposure.*

2 Paling, Emma. (June 18, 2020). "Canada's Coronavirus Tracing App will be Piloted in Ontario." *Huffington Post.*

3 Fahey, Robert & Hino, Airo. (2020). "COVID-19, Digital Privacy, and the Social Limits on Data-Focused Public Health Responses." *International Journal of Information Management,* 55.

4 Amnesty International. (April 3, 2020). "COVID-19, surveillance and the threat to your rights."; Davis, Jessica. (April 20, 2020). "ACLU, Scientists Urge Privacy Focus for COVID-19 Tracing Technology." *Health IT Security.*

5 Gray, Stacey. (March 25, 2020). "A Closer Look at Location Data: Privacy and Pandemics." *Future of Privacy Forum.*

6 Cofone, Ignacio & Robertson, Adriana. (2018). "Consumer Privacy in a Behavioral World." *Hastings Law Journal*, 69, 1471-1508, p. 1496.

7 Semeniuk, Ivan. (June 18, 2020). "Ottawa promotes contact tracing app for Canadians in fight against the spread of COVID-19." *The Globe and Mail.*; Office of the Prime Minister. (October 5, 2020). *Canada's COVID-19 exposure notification app now available in Quebec.*; Office of the Ontario Premier. (October 15, 2020). *Millions Across Canada Now Using Made-in-Ontario COVID Alert App.*

8 *CBC News*. (October 30, 2020). "Canada's COVID-19 Alert app updated to include more precise exposure information."

9 Office of the Prime Minister. (July 31, 2020). Prime Minister's remarks on COVID-19 measures and the launch of the COVID Alert national application.; *The Canadian Press*. (July 31, 2020). "Canadians can now download new COVID-19 exposure-alert smartphone app." *The Canadian Press.*

10 Government of Ontario. (2020). *What to do if you've been exposed to COVID-19.*

11 *Privacy International.* (March 31, 2020). "Bluetooth tracking and Covid-19: a tech primer."

12 Ibid.

13 Daigle, Thomas. (August 13, 2020). "Misconceptions persist about effectiveness and privacy of Canada's COVID Alert app." *CBC News.*

14 Government of Canada. (October 2020). *COVID Alert COVID-19 Exposure Notification Application Privacy Assessment*.

15 Google. (September 2020). "Exposure Notification: Frequently Asked Questions," p. 2.

16 Google. (2020). "Exposure Notifications: Using technology to help public health authorities fight COVID-19."

17 Baumgärtner, Lars et al. (2020). "Mind the GAP: Security & Privacy Risks of Contact Tracing Apps." *ArXiv*, p. 4-5.

18 Tanguay-Renaud, François et al. (2020). "Test, Trace, and Isolate: Covid-19 and the Canadian Constitution." *Osgoode Legal Studies Research Paper,* p. 6.

19 MacKinnon, Bobbi-Jean. (May 7, 2020). "Canada's privacy commissioners offer guidance on COVID-19 contact-tracing Apps." *CBC News.*

20 Ibid.

21 European Data Protection Board. (April 21, 2020). "Guidelines 04/2020 on the use of location data and contact tracing tools in the context of the COVID-19 outbreak," paras 24 and 27.

22 Fahey & Hino, *supra* note 3.

23 Ruths, Derek. (June 18, 2020). "Canada's proposed contact-tracing app takes the right approach on privacy." *The Globe and Mail.*

24 Fahey & Hino, *supra* note 3.

25 Ruths, *supra* note 23.

26 Cofone, Ignacio & Robertson, Adriana. (2018). "Privacy Harms." *Hastings Law Journal*, 69, 1039-1098, p. 1044-1056.

27 Cofone, Ignacio. (2020). "Nothing to Hide, but Something to Lose." *Toronto Law Journal*, 70(1), 1-45, p. 4, 30-32.

28 Tanguay-Renaud, François et al., *supra* note 18.

29 Cyphers, Bennett and Gebhart, Gennie. (April 28, 2020). "Apple and Google's COVID-19 Exposure Notification API: Questions and Answers." *Electronic Frontier Foundation*.

30 Solove, Daniel J. (2013). "Privacy Self-Management and the Consent Dilemma." *Harvard Law Review*, 126(7), 1880-1903, p. 1892.

31 Lanzing, Marjolein. (forthcoming 2020). "Contact Tracing Apps: An Ethical Roadmap." *Ethics in Information Technology,* 1-4, p. 3.

32 O'Neill, Patrick Howell. (June 5, 2020). "No, coronavirus apps don't need 60% adoption to be effective." *MIT Technology Review.*

33 Ferretti, Luca et al. (2020). "Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing." *Science*, 368 (6491) 1-7.

34 Farronato, Chiara et al. (July 15, 2020). "How to Get People to Actually Use Contact-Tracing Apps." *Harvard Business Review.*

35 Kleinman, Robert and Merkel, Colin. (2020). "Digital contact tracing for COVID- 19." *CMAJ,* 1-4.

36 Ruihley, Josh et al. (July 31,2020). "Continuously improving COVID Alert." *Canadian Digital Services* (blog).

37 Baumgärtner et al., *supra* note 17.

38 Rosner, Gilad. (2020). "De-Identification as Public Policy." *Journal of Data Protection and Privacy.* 3(3), 1-18, p. 1.

39 Rubenstein, Ira & Hartzog, Woodrow. (2016). "Anonymization and Risk." *Washington Law Review*, 91(2) 703-760, p. 704-730.

40 Barocas, Solon & Nissenbaum, Helen. (2014). "Big Data's end turn around procedural privacy protections." *Communications of the ACM*, 57(11), 31-33.

41 Stanley, Jay & Jennifer Stisa Granick. (April 8, 2020). "ACLU White Paper: The Limits of Location Tracking in an Epidemic." *ACLU*.

42 Greenberg, Andy. (April 17, 2020). "Does Covid-19 Contact Tracing Pose a Privacy Risk? Your Questions, Answered." *Wired*.

43 Thompson, Elizabeth. (August 5, 2020). "COVID Alert app could result in some people being ID'd." *CBC News.*

44 Barocas, Solon, & Levy, Karen. (2020). "Privacy dependencies." *Washington Law Review* 95, 555–616.

45 Kim, Pauline T. (2017). "Data-Driven Discrimination at Work." *William & Mary Law Review*, 58(3), 857–936, p. 869–92; Barocas, Solon & Selbst, Andrew D. (2016). "Big Data's Disparate Impact." *California Law Review*, 104, 671–732, p. 673–76.

46 See e.g. Jones et al. (2018). "Toward an Ethically Founded Framework for the Use of Mobile Phone Call Detail Records in Health Research." *JMIR Mhealth & Uhealth*, 7(3), 1-6, p. 6.

47 See e.g., Finn, Rachel L., Wright, David & Friedwald, Michael. (2013) "Seven Types of Privacy" in Serge Gutwirth et al, eds, *European Data Protection: Coming of Age* (Springer), p. 3-32.

48 Cofone, Ignacio & Strandburg, Katherine. (2019). "Strategic Games and Algorithmic Secrecy." *McGill Law Journal*, 64(4) 623–663, p. 634–636.

49 Greenberg, *supra* note 42.

50 Tanguay-Renaud et al., *supra* note 18, p. 3.

51 Simko, Lucy, et al. (2020). "COVID-19 Contact Tracing and Privacy: Studying Opinion and Preferences." ArXiv, 1-32.

52 Ibid, p. 6.

53 O'Neil, Cathy. (2013). *On Being a Data Skeptic.* (O'Reilly Media, Inc), p. 1-19.

54 Cyphers & Gebhart, *supra* 29.

55 Cofone, Ignacio. (2019). "Algorithmic Discrimination is an Information Problem." *Hastings Law Journal*, 70(2), 1389-1444, p. 1389, 1406.

56 Madden, Mary et al. (2017). "Privacy, Poverty, and Big Data." *Washington University Law Review*, 95(1), 53-125.

57 Davis, Sara. (April 29, 2020). "Contact Tracing Apps: Extra Risks for Women and Marginalized Groups." *Health and Human Rights Journal.*

58 Ibid.

59 Cyphers & Gebhart, *supra* note 29.

60 See, e.g., Frankle, Jonathan & Garvie, Clare. (April 7, 2016). "Facial-Recognition Software Might Have a Racial Bias Problem." *The Atlantic.*

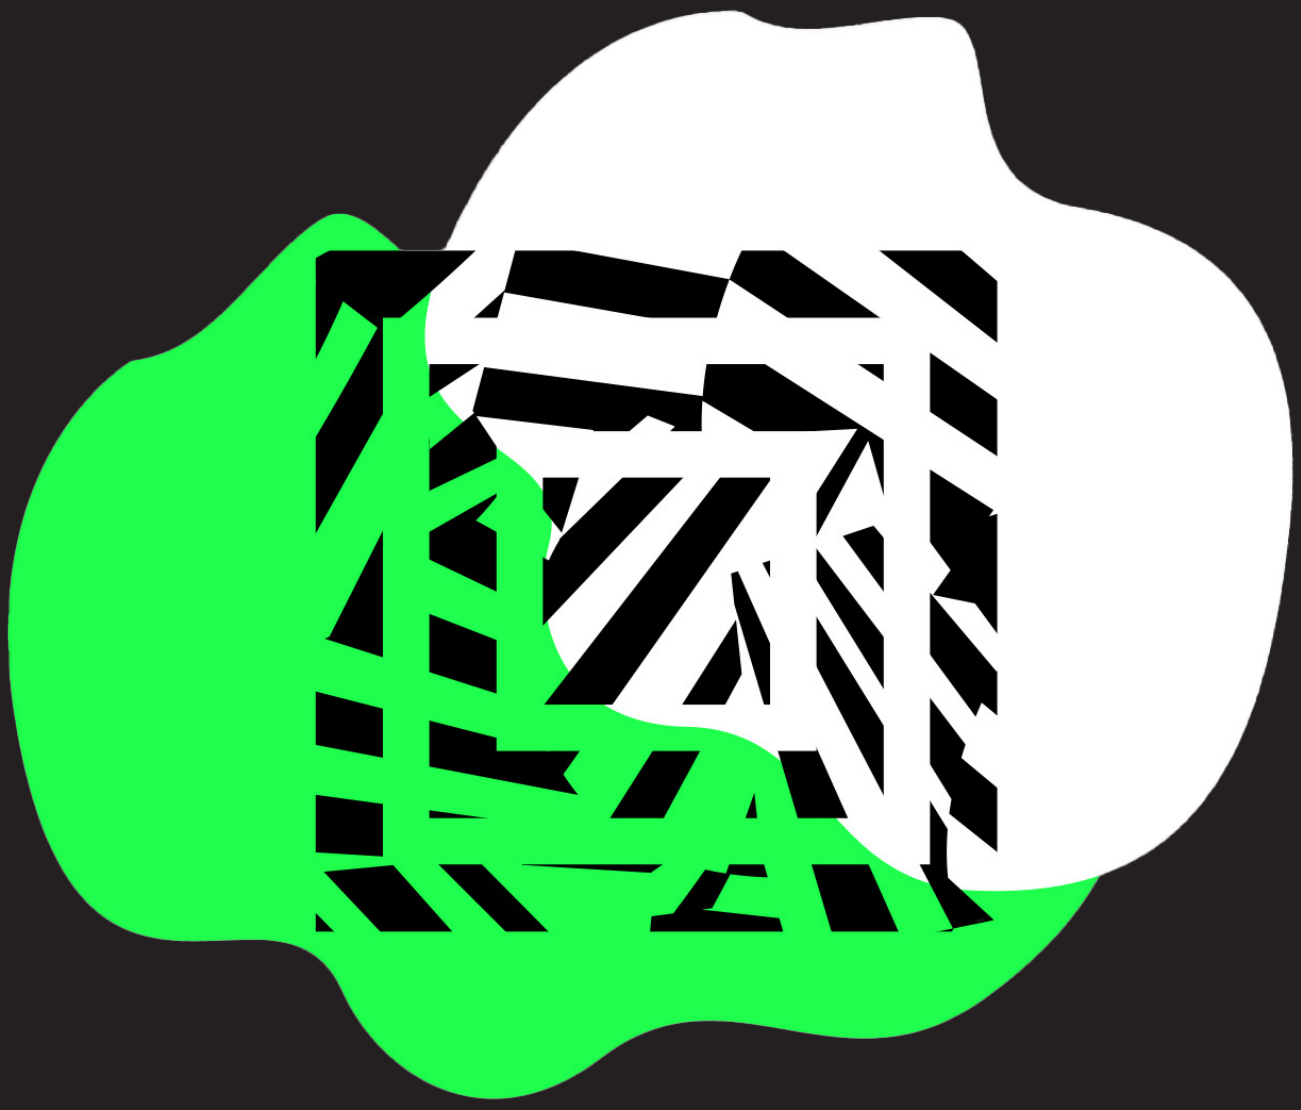

Designed by Codi Hauka and Yasmeen Safaie

Centre for MEDIA, TECHNOLOGY and DEMOCRACY
McGill UNIVERSITY
MAX BELL SCHOOL of PUBLIC POLICY