BROOKINGS TECH

Tomorrow's tech policy conversations today

About TechStream | Stay Informed

Contact-tracing apps are not a solution to the COVID-19 crisis

April 27, 2020 | Ashkan Soltani, Ryan Calo, and Carl Bergstrom

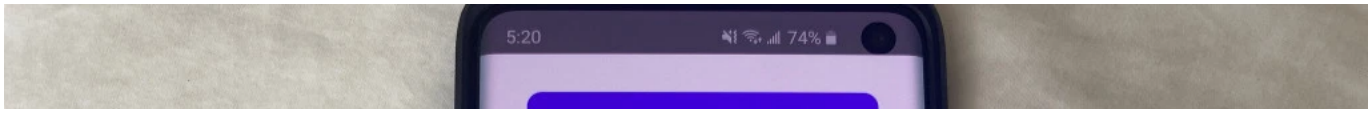

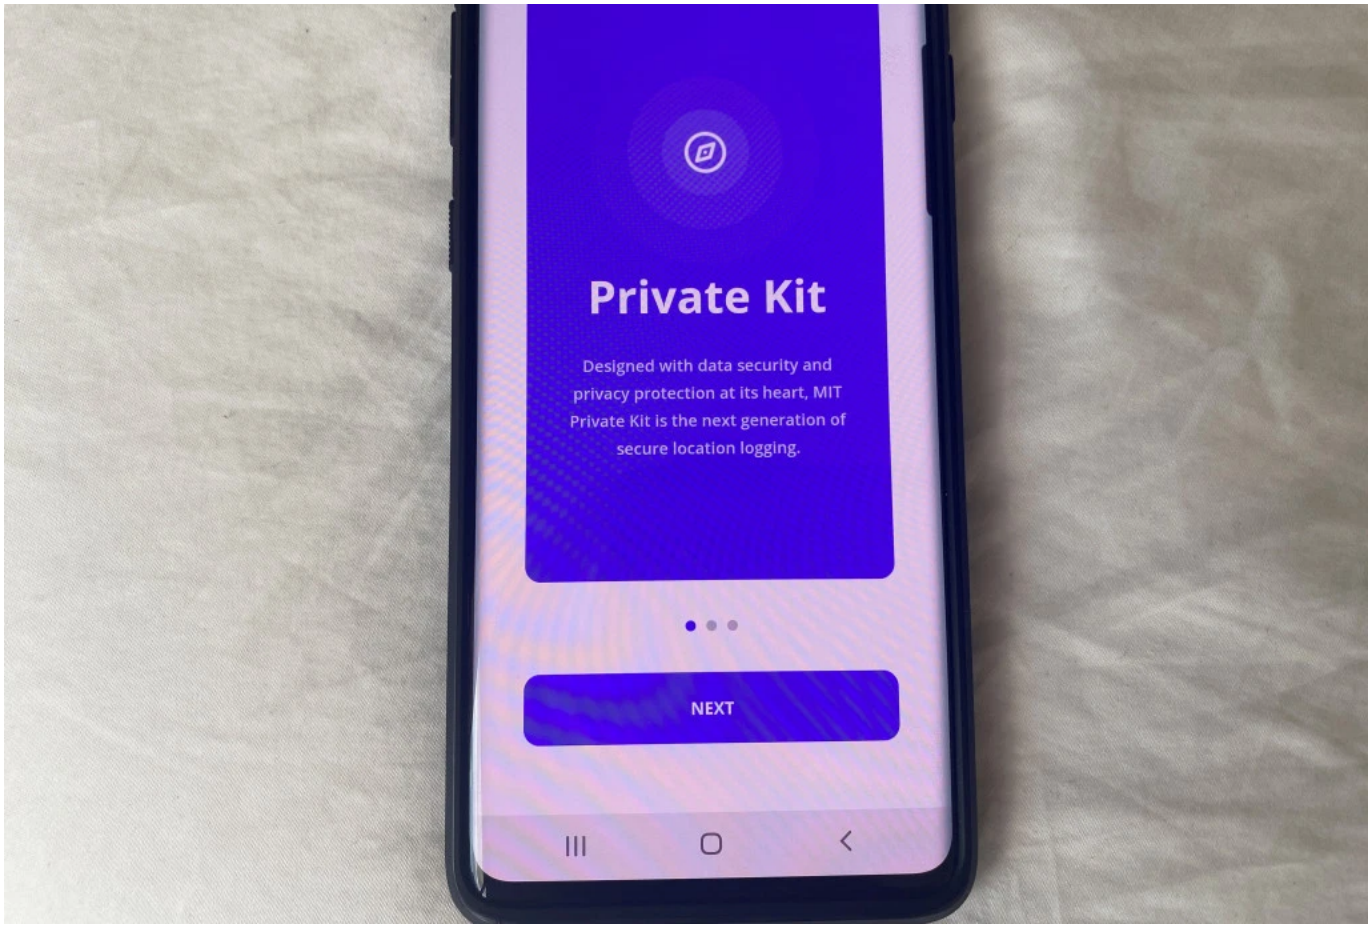

The unprecedented threat from the novel coronavirus has confined many Americans to their homes, distancing them from one another at great cost to local economies and personal well-being. Meanwhile the pressure grows on American institutions to do something—anything—about the pandemic.

Encouraged by the White House, much of that pressure to act has focused on Silicon Valley and the tech industry, which has responded with a fragile digital solution. Tech companies and engineering departments at major universities are pinning their hopes of returning Americans to work and play on the promise of smartphone apps. Coronavirus? There's an app for that.

We are concerned by this rising enthusiasm for automated technology as a centerpiece of infection control. Between us, we hold extensive expertise in technology, law and policy, and epidemiology. We have serious doubts that voluntary, anonymous contact tracing through smartphone apps—as Apple, Google, and faculty at a number of academic institutions all propose—can free Americans of the terrible choice between staying home or risking exposure. We worry that contact-tracing apps will serve as vehicles for abuse and disinformation, while providing a false sense of security to justify reopening local

and national economies well before it is safe to do so. Our recommendations are aimed at reducing the harm of a technological intervention that seems increasingly inevitable.

We have no doubts that the developers of contact-tracing apps and related technologies are well-intentioned. But we urge the developers of these systems to step up and acknowledge the limitations of those technologies before they are widely adopted. Health agencies and policymakers should not over-rely on these apps and, regardless, should make clear rules to head off the threat to privacy, equity, and liberty by imposing appropriate safeguards.

Proposals to combat coronavirus using smartphones largely focus on facilitating the process of "contact tracing." Contact tracing involves working backward from infected cases to identify people who may have been exposed to disease, so that they can be tested, isolated, and—when possible—treated. Traditional contact tracing is a labor-intensive process of interviews and detective work. Some countries such as Singapore, South Korea and Israel have enlisted technology, including mobile apps, to facilitate contact tracing of coronavirus cases, and this idea is now catching on in the United States. North Dakota and Utah have released voluntary contact-tracing apps that rely on tracking users' location as they move about, and the consulting firm PwC has begun promoting a contact-tracing tool to permit employers to screen which employees can return to work. Several American technology companies and institutions of higher learning are developing the infrastructure that would permit automated contact tracing of a sort, while also avoiding certain privacy concerns.

Contact tracing can be an important component of an epidemic response especially when the prevalence of infection is low. Such efforts are most effective where testing is rapid and widely available and when infections are relatively rare —conditions that are currently unusual in the United States. Ideally, manual contact tracing by trained professionals can help identify candidates for testing and quarantine to help contain the spread of coronavirus.

The lure of automating the painstaking process of contact tracing is apparent. But to date, no one has demonstrated that it's possible to do so reliably despite

numerous concurrent attempts. Apps that notify participants of disclosure could, on the margins and in the right conditions, help direct testing resources to those at higher risk. Anything else strikes us as implausible at best, and dangerous at worst.

In response to increased pressure from the Trump administration on technology platforms to share data, Apple and Google have proposed an application programming interface (or "API") for conducting contact tracing using mobile phones, which they describe as a system to provide "exposure notification" to users once they've been diagnosed or self-report as infected. The Apple-Google API supports the specific functionality of warning participants if their phone has been near the phone of a person who reported being COVID-19 positive. To be clear, the companies are not planning to develop an app themselves, which would require addressing some of the more challenging issues around how to verify that a user has been infected and what policies to suggest when individuals are alerted to being "in contact" with an infected individual. Ultimately, they have left it up to public health officials, or whoever else develops the apps, to determine their functionality and uses—subject, of course, to the constraints of the platform.

We and many others have pointed out a host of pitfalls for voluntary, self-reported coronavirus apps of the kind Apple, Google, and others contemplate. First, app notifications of contact with COVID-19 are likely to be simultaneously both over- and under-inclusive. Experts in several disciplines have shown why mobile phones and their sensors make for imperfect proxies for coronavirus exposure. False positives (reports of exposure when none existed) can arise easily. Individuals may be flagged as having contacted one another despite very low possibility of transmission—such as when the individuals are separated by walls porous enough for a Bluetooth signal to penetrate. Nor do the systems account for when individuals take precautions, such as the use of personal protective equipment, in their interactions with others.

Even among true contact events, most will not lead to transmission. Studies suggest that people have on average about a dozen close contacts a day—incidents involving direct touch or a one-on-one conversation—yet even in the

absence of social distancing measures the average infected person transmits to only 2 or 3 other people throughout the entire course of the disease. Fleeting

interactions, such as crossing paths in the grocery store, will be substantially more common and substantially less likely to cause transmission. If the apps flag these lower-risk encounters as well, they will cast a wide net when reporting exposure. If they do not, they will miss a substantive fraction of transmission events. Because most exposures flagged by the apps will not lead to infection, many users will be instructed to self-quarantine even when they have not been infected. A person may put up with this once or twice, but after a few false alarms and the ensuing inconvenience of protracted self-isolation, we expect many will start to disregard the warnings. Of course this is a problem with conventional contact tracing as well, but it can be managed with effective direct communication between the contact tracer and the suspected contact.

At least as problematic is the issue of false negatives—instances where these apps will fail to flag individuals as potentially at risk even when they've encountered someone with the virus. Smartphone penetration in the United States remains at about 81 percent—meaning that even if we had 100 percent installation of these apps (which is extremely unlikely without mandatory policies in place), we would still only see a fraction of the total exposure events (65 percent according to Metcalf's Law). Furthermore, people don't always have their phones on them. Imagine the delivery person who leaves her phone in the car. Or consider that the coronavirus can be transmitted via the surfaces on which it lingers long after a person and their phone has left the area. The people in the highest risk groups—the aging or under-resourced—are perhaps least likely to download the app while needing safety most. Others may download the app but fail to report a positive status—out of fear, because they are never tested, or because they are among the significant percentage of carriers who are asymptomatic.

Contact-tracing apps therefore cannot offer assurance that going out is safe, just because no disease has been reported in the vicinity. Ultimately, contact tracing is a public health intervention, not an individual health one. It can reduce the spread of disease through the population, but does not confer direct protection on any individual. This creates incentive problems that need careful thought: What is in it for the user who will sometimes be instructed to miss work and avoid socializing, but does not derive immediate benefits from the system?

Some of the contact-tracing frameworks have been designed with security and privacy in mind, to some degree. The Apple-Google proposal, for example, stores

the information about what "contacts" the device has made on each users' device, rather than reporting that information to a central server as is the case with some of the other approaches. This "decentralized" architecture isn't completely free of privacy and security concerns, however, and actually opens apps based on these APIs to new and different classes of privacy and security vulnerabilities. For example, because these contact-tracing systems reveal health status in connection with a unique (if rotating) identifier, it is possible to correlate infected people with their pictures using a stationary camera connected to a Bluetooth device in a public place.

And finally, the issue of malicious use is paramount—particularly given this current climate of disinformation, astroturfing, and political manipulation. Imagine an unscrupulous political operative who wanted to dampen voting participation in a given district, or a desperate business owner who wanted to stifle competition. Either could falsely report incidences of coronavirus without much fear of repercussion. Trolls could sow chaos for the malicious pleasure of it. Protesters could trigger panic as a form of civil disobedience. A foreign intelligence operation could shut down an entire city by falsely reporting COVID-19 infections in every neighborhood. There are a great many vulnerabilities underlying this platform that have still yet to be explored.

Though technologists at Apple, Google, and a number of academic institutions have given some thought in their planning documents to the possibility that their tools could be exploited and abused, they need to be much more candid about the limitations of the technology—including the fact that these approaches should never be used in isolation, if they are used at all. Like thermometers, tires, and many other products that operate safely only within a specific range, these apps should come with a warning about their many points of failure.

There is also a very real danger that these voluntary surveillance technologies will effectively become compulsory for any public and social engagement. Employers, retailers, or even policymakers can require that consumers display the results of

their app before they are permitted to enter a grocery store, return back to work, or use public services—is as slowly becoming the norm in China, Hong Kong, and

even being explored for visitors to Hawaii.

Taken with the false positive and "griefing" (intentionally crying wolf) issues outlined above, there is a real risk that these mobile-based apps can turn unaffected individuals into social pariahs, restricted from accessing public and private spaces or participating in social and economic activities. The likelihood that this will have a disparate impact on those already hardest hit by the pandemic is also high. Individuals living in densely populated neighborhoods and apartment buildings—characteristics that are also correlated to non-white and lower income communities—are likelier to experience incidences of false positives due their close proximity to one another.

Therefore, we urge developers of contact-tracing apps, as well the companies enabling their development, to be candid about the limitations and implications of the technology. To be ethical stewards of these new public health tools, they must also provide explicit guidelines and "best practice" recommendations for the development of the apps. These should include recommendations for how back-end systems should be secured and how long data should be retained, criteria for what public health entities can qualify to use these technologies, and explicit app store policies for what additional information, such as GPS or government ID numbers, can be collected. They should adopt commonly accepted practices such as security auditing, bug bounties, and abusability testing to identify vulnerabilities and unintended consequences of a potentially global new technology. Finally, app creators—as well as the platforms that enable these applications—should make explicit commitments for when these apps and their underlying APIs will be sunsetted.

There is also a role for law and official policy. If we are to use technology to combat coronavirus, it is critical that we do so with adequate safeguards in place. Here we mean traditional safeguards, such as judicial oversight and sunset provisions that guard against mission creep or limitations on secondary use and data retention that protect consumer privacy. We agree with our colleagues at the

Civil Liberties Oversight Board that coronavirus surveillance should learn from the lessons of 9/11. But we also see a role of law and policy in policing against an

all too plausible dystopia that technological solutions could enable.

Lawmakers, for their part, must be proactive and rapidly impose safeguards with respect to the privacy of data, while protecting those communities who can be— and historically have been— harmed by the collection and exploitation of personal data. Protections need to be put in place to expressly prohibit economic and social discrimination on the basis of information and technology designed to address the pandemic. For example, academics in the United Kingdom have proposed model legislation to prevent compulsory or coerced use of these untested systems to prevent people from going back to work, school, or accessing public resources. The prospect of surveillance during this crisis only serves to reveal how few safeguards exist to consumer privacy, especially at the federal level.

At the end of the day, no clever technology—standing alone—is going to get us out of this unprecedented threat to health and economic stability. At best, the most visible technical solutions will do more than help on the margin. At a minimum, it is the obligation of their designers to ensure they do no harm.

Ashkan Soltani *is an independent researcher and technologist specializing in privacy, security, and behavioral economics. He was previously a senior advisor to the U.S. Chief Technology Officer, the chief technologist for the Federal Trade Commission, and a contributor to the Washington Post team that in 2014 won a Pulitzer Prize for its coverage of national-security issues.*

Ryan Calo *is a professor of law at the University of Washington, with courtesy appointments in computer science and information science and the co-founder of two interdisciplinary research initiatives.*

Carl Bergstrom *is a professor of biology at the University of Washington with extensive experience in the epidemiology of emerging infectious diseases, which he is integrating into ongoing research on spread of disinformation through social and traditional media channels during the SARS-CoV-2 pandemic.*