

Privacy, Adoption, and Truthful Reporting: A Simple Theory of Contact Tracing Applications

Yves-Alexandre de Montjoye, Tarun Ramadorai,
Tommaso Valletti, and Ansgar Walther

Imperial College London

November 21, 2020

Abstract

This short paper analyses the trade-offs associated with the deployment of contact tracing applications to support policy responses in a pandemic. In many jurisdictions, the government cannot force individuals to adopt such applications. We therefore analyze a simple model that highlights the importance of individuals' incentives to voluntarily adopt a reporting application and reveal their infection status to the government who can then undertake contact monitoring. We discuss the consequences of various policy options, such as security, communication and anonymisation policies, in terms of the size and representativeness of the sample of infection data that contract tracing applications generate.

1 Introduction

A strategy that governments and public health authorities around the world have pursued to combat the coronavirus pandemic is to develop and seek widespread adoption of smartphone apps that can aid contact monitoring and tracing. Contact tracing applications are viewed as an important part of the toolkit to deal with a public health emergency, but their usefulness relies on their adoption and use by a representative sample of individuals. If the apps are broadly adopted and utilized, they can provide readily available, accurate, and granular information about infection status, which in turn can enable governments to swiftly enforce localized lockdowns or allow targeted groups of individuals to self-quarantine and break the chain of propagation. This in turn can help to halt the broader spread of the disease, whilst limiting broader damage to the economy.

Government efforts have been focused on realizing the potential benefits to the broader community of the information gathered from these apps, such as more efficiently targeted and implemented lockdown or quarantining strategies. On the other hand, an intense debate has focused on the potential privacy and security consequences for individuals from adopting these apps. Apart from the possible privacy consequences of detailed data gathering on individual movements, there are also issues associated with the vulnerability of different technological solutions to attacks from malicious actors,¹ and on the direct consequences of reporting individual infection status on the liberty and freedom of movement of the adopter and their close contacts. In some parts of the world, states can attempt to coerce citizens to adopt such apps and to force them to share their personal data with public health authorities. However, such forced adoption is antithetical to the basic nature of most democracies, provoking important discussions about rights and trade-offs.² Moreover, even coerced compliance in non-democratic societies is likely

¹There is considerable debate on the best protocols to be utilized, with some consensus developing on bluetooth protocols as the most promising. The details of implementation are critical, and numerous solutions have emerged including decentralized solutions (such as DP3T, Gapple, and MIT PACT), centralized solutions (e.g., NHS, ROBERT), and third-way (e.g., DESIRE). Imperial College's 8 privacy questions to evaluate contact tracing apps are summarized in this blogpost.

²Indeed the legislative framework in many countries has been designed to ensure contact tracing apps are and will stay a purely voluntary decision. This includes very strong penalties for trying to make it mandatory (recently, for example, the Belgian Data Protection Authority went as far as asking for criminal sanctions for mandatory app downloads).

to face widespread evasion and reluctance to participate. In all cases, therefore, understanding the incentives of individuals that lead to their decisions of whether and how to use tracking apps seems critical.

In this short note, we provide a simple framework to document the trade-offs associated with these questions, and to study individual decisions to install and adopt apps and truthfully report. Our approach uses basic theoretical frameworks from economics to outline some trade-offs associated with the rollout of tracing apps to acquire information to support policy responses in the pandemic, the representativeness of the data acquired from these apps, and the net benefits perceived by users in engaging in downloads.

To fix ideas, consider the example of young and healthy individuals, who receive a wage only if they turn up to work, with low financial savings and limited access to government furlough or other insurance schemes. Voluntary adoption of, and compliance with, tracing app downloads is relatively costly for such individuals. Since they belong to a relatively low health risk category, even conditional on contracting the virus, they may face mild symptoms. However, if they were to report their infection status, the resulting impact on their income from compliance with a fourteen day quarantine regime is substantially negative. These are direct costs of compliance, in addition to any costs arising from their preferences for privacy, or the possibility of being subject to malicious technological attacks because of the app download. That said, the social value arising from their data provision is potentially quite high, and comparing private costs to public benefits engenders a classic free-rider problem.

We develop these ideas further in our simple model below, casting the problem in more general terms to determine the factors leading to the individual's decision to contribute personal data, and the factors permitting the government to break the chain of propagation and accurately assess infection rates based on individual decisions to adopt and report. While our exercise is less detailed than those in the combined epidemiological and economic models that are the focus of a fast-growing literature (see, e.g., Atkeson 2020, Berger et al. 2020, Eichenbaum et al. 2020, and Stock 2020), the value added of our approach comes from the ability to outline and formalize the trade-offs facing public policy in analyzing an important aspect of pandemic response. The basic dilemma that we study is that individual preferences for pri-

vacy and individual consequences from supplying data may differ substantially from the social benefits accruing from decisions to adopt and truthfully report. This wedge is a more general problem in the data economy, and studying this problem also helps to shed light on the broader study of data privacy, a fast-growing area of study (see, e.g., Varian, 2009, 2010, Goldfarb and Tucker, 2011, 2012, Acquisti et al., 2015, and Ramadorai et al., 2019).

2 A simple model of adoption

Each individual in a population observes a signal $x \in [0, 1]$, which reveals their likelihood of being infected. Let $F(x)$ be the (cumulative) posterior distribution of perceived infection risk in the population, conditional on individual information received. Individuals can adopt a reporting technology which reveals verifiable information about their type to the government. We write $a \in \{0, 1\}$ for the decision to adopt this reporting technology ($a = 1$) or not ($a = 0$). In addition to their infection probability x , each individual observes their private net benefit $b \in [\underline{b}, \bar{b}]$ that they would derive from adopting the technology. Let $H(b|x)$ be the distribution of the net benefit b among people with infection risk x .

Since it is a net benefit, the parameter b incorporates costs as well as advantages of adoption. On the cost side, b will be lower if there are potential restrictions on infected persons' freedom of movement or ability to work. Another possible cost is that of personal data being accessed by third parties as a result of a leak or a malevolent attack on the data sharing technology (see, e.g., Radaelli, et al., 2018, Vaudenay, 2020). On the benefit side, b will be higher if individuals and those in their community benefit from personalized care if they reveal their vulnerability to the government. Another possible component of the benefit can arise from social preferences, for example, from the “warm glow” utility associated with contributing to the common goal of fighting the epidemic (e.g., Morewedge et al., 2009, Bordalo et al., 2012). Importantly, we allow the distribution of b to depend on x . This dependence is natural in a world where the reporting technology passes individual, non-anonymised information to the government—the government is more likely to impose restrictions on, or provide special care to, an individual who has reported a higher likelihood of being infected.

The government employs information gathered using the reporting technology to provide public health services that can be accessed by both infected and non-infected individuals. We assume that once the government has used this information optimally, each infected individual derives utility:

$$V(P, D, x)$$

from health services (which allows for potentially different benefits conditional on infection likelihood).

Here,

$$P = Pr [a = 1]$$

denotes the population-wide adoption rate, and

$$D = E [x|a = 1] - E [x]$$

is a measure of the deviation in infection rates between a representative sample provided by adopting individuals and the true estimate of the average infection rate in the population. That is, if $D = 0$, then the average infection rate among reporters is a good estimate of the average infection rate in the population. If $D > 0$, then the government's sample is positively selected and over-represents infection rates, and it is negatively selected if $D < 0$. We further assume that $\frac{\partial V}{\partial P} > 0$, and $\frac{\partial V}{\partial D} \stackrel{\text{sign}}{=} -D$, so that large, representative samples allow the government to choose better policies (i.e., value V is maximized when $D = 0$). Since each individual is infected with probability x , she derives utility $V(P, D, x)$ from public health services.

An additional natural assumption, although it is not required for the results we derive below, is that $\partial V(P, D, x)/\partial x > 0$, so that individuals with high infection risk benefit more from public health policy.

We make a few comments on this very simple setting. First, the revelation of individuals' "type" is not as rich as reality, where multiple individual revelations are combined by a central agency to trace out a network of contacts and likely infections. Our model is a simplification, in that participation conveys "types". Despite this simplicity, the model captures the idea that

giving away information to a central agency can be costly and needs to be incentivized. These incentives are the focus of our analysis, and we believe that they would operate similarly in a more fully specified model of contact networks. While such a more specified model would render the associated costs and benefits more concrete, our approach provides some essential insights in a more reduced-form fashion.

Second, we assume that the observed signal x can be verifiably disclosed to the government. We make this assumption to simplify the exposition. In reality, one might expect partial verifiability, in the sense that some aspects of individuals' infection risk are easier to verify than others. For example, verification of individuals' reported test results is relatively straightforward, but it is more difficult to verify an app user's report of their contacts with others.

Third, the observed signal x does not necessarily have to come from the app (e.g., an individual will observe his/her own symptoms). However, the signal may also come via the app (e.g., a person is notified of their proximity to someone who tested positive). In either case, the question is whether to act upon the signal (e.g., in both cases, to take a test and report the results to the authorities).

Finally, we note that $V(P, D, x)$ in our specification is an "indirect" utility function, in the sense that it captures the expected utility that individuals enjoy once the government has chosen the optimal health policy. Consequently, while $V(P, D, x)$ does not depend directly on policy choices in our formalism, this dependence is implicit, since we assume that the government chooses the best possible policy from a social perspective, using available information.

2.1 Adoption decisions

In this simple model, the total utility of an individual with type (x, b) who takes adoption decision a is:

$$V(P, D, x) + ba$$

Thus she optimally adopts the technology if and only if

$$b \geq 0$$

Although mathematically trivial, this contains an important point: Individual adoption decisions will be driven by *private* net benefits and social preferences, not by individual valuations of the *public* health benefit $V(P, D, x)$.

One might intuitively expect, for example, that individuals with high infection risk x are more likely to adopt the technology because they benefit most from health services. However, this intuition is misleading, because the public health benefit V depends on *aggregate* adoption decisions via the sample size and representativeness parameters P and D , but is not affected by the decisions of one (infinitesimally) small individual.

An extreme implication of this result is that *nobody* will adopt the technology if there are only private costs associated with adoption, regardless of how large the public benefits are. This is an instance of the classic free-rider problem in public goods provision; the main difference from the textbook case is that free-riding occurs via a lack of data provision, rather than a lack of monetary contributions. This intuition may have been behind the experience in France, where Covid-19 app adoption has been sluggish, hampering public health efforts.³

2.2 Sampling properties

Our analysis in the previous section implies that an individual with infection risk x adopts the technology with probability:

$$\begin{aligned} Pr[a = 1] &= Pr[b \geq 0|x] = 1 - H(0|x) \\ &\equiv p(x) \end{aligned}$$

Thus the total rate of adoption in the population is

$$\begin{aligned} P &= Pr[a = 1] = \int_0^1 p(x)dF(x) \\ &= E[p(x)] \end{aligned}$$

³“French Contact-Tracing App Struggles with Slow Adoption. It Isn’t Alone”, WSJ, June 23, available at <https://www.wsj.com/articles/french-contact-tracing-app-struggles-with-slow-adoption-it-isnt-alone-11592928266>

Moreover, to derive the representativeness metric, note that

$$\begin{aligned} E[x|a = 1] &= \frac{1}{P} \int_0^1 xp(x)dF(x) \\ &= \frac{E[x \cdot p(x)]}{P} \\ &= E[x] + \frac{Cov[x, p(x)]}{P} \end{aligned}$$

Hence we have

$$D = \frac{Cov[x, p(x)]}{P}$$

Two insights emerge. First, the government's sample is perfectly representative (i.e., $D = 0$) if and only if the infection risk and the adoption propensity are uncorrelated (i.e., $Cov[x, p(x)] = 0$). Moreover, the government's sample is positively (negatively) selected if and only if the infection risk x correlates positively (negatively) with the propensity $p(x)$ to adopt the technology. For example, if $p'(x) > 0$, then $D > 0$ and if $p'(x) < 0$, then $D < 0$.

Second, adoption rates $P = \Pr[a = 1] = E[p(x)]$ and representativeness D are nonlinearly related to one another (i.e., they have an interactive relationship). Consider a case in which the sample is not representative because it is positively selected (i.e., $D > 0$). In this case, increasing the adoption rate (upping P) will also increase representativeness (i.e., shrink D towards zero).

To understand the intuition for this result, consider a very simple case in which the net benefit from adoption is perfectly correlated with x , so that $b = \alpha + \beta x$. If $\beta > 0$, then the adopters are all individuals with $x \geq -\frac{\alpha}{\beta}$ (since adoption requires $b > 0$), i.e., the high-risk group in the population. If $\beta < 0$, by contrast, then the adopters are those with infection risk $x \leq -\frac{\alpha}{\beta}$, i.e., the low-risk group in the population. In this situation, if the baseline benefit α increases, this increases participation P (since both conditions become more slack), and the sample also becomes more representative in both cases.

3 Policies, privacy, and adoption

We can use these derivations to discuss what types of policies might increase adoption rates. Clearly, any policy that increases perceived net benefits b will increase participation by raising the participation probabilities $p(x)$. For example, one natural policy to increase participation is to enhance the *security* of the data collection exercise. As we argued above, net benefits b are lower when individuals expect the reporting technology to be vulnerable to hacks or attacks. Another insight from the economic analysis of this problem is that one should target security policies at *marginal* adopters with benefits just below zero. It is of no use, in terms of overall participation, to put in place measures that reassure extremely privacy-conscious individuals, because these individuals are very unlikely to switch their behavior to adoption since their perceived net benefits are likely to be extremely low. This also raises an additional issue, which is that we assume that the distribution of perceived benefits is uncorrelated with the infection rate. This is a potentially testable hypothesis, and pending data on this issue, a possibility is to extend this analysis to account for such a case. Finally, if one believes that part of the net benefit b is driven by social preferences, then another policy to increase adoption is communication and positive messaging. Encouraging the population to do their social duty can then have the effect of shifting the distribution of b upwards, thereby increasing participation rates.

3.1 Increasing representativeness using anonymity guarantees

One important lesson from the above analysis is that participation breeds representativeness: Any increase in aggregate participation P shrinks the deviation D from a representative sample towards zero. Therefore, any of the above participation-based policies are also useful to generate a representative sample. However, in cases where it is not feasible to increase participation further, it may be useful instead to work on directly reducing deviations from a representative sample. Our analysis shows that the key step in this exercise is to decorrelate infection risk x from the perceived costs and benefits of adoption which drive $p(x)$. One obvious candidate for a decorrelation policy is an anonymity guarantee. If it can be guaranteed that no individual is identifiable after revealing their infection status, then the private costs and benefits associated

with restrictions and personalised care conditional on infection status disappear, breaking the link between x and $p(x)$. Hence, a policy which removes considerations of penalties or benefits associated with infection status from individuals' calculations by truly anonymizing this information will also help to increase representativeness.

Of course, while it is clear that this type of decorrelation policy has the potential benefit of increasing the representativeness of the government's dataset, it must be noted that such anonymity guarantees also come at substantial costs. In an epidemic, anonymity severely limits the government's ability to crack down on local outbreaks, precisely because it removes the ability to impose targeted restrictions on individuals with a high likelihood of infection, or indeed their close contacts, as reaching out to them could violate any anonymity guarantees. A solution might then be to impose a less precisely targeted policy, i.e., a broader lockdown even in the presence of more precise information. The excess costs of such a broader lockdown would then be traded off against the greater representativeness of data—acquired through higher adoption rates arising from the provided anonymity.

3.2 Improving statistical technology

We note that, even with a sample that is not representative, one can estimate average infection rates accurately using inverse propensity score weighting. If one takes the average of reported x divided by adoption propensity scores $p(x)$, one obtains:

$$E \left[\frac{x}{p(x)} | a = 1 \right] = E[x]$$

This is obviously not feasible as a direct estimator, because the infection risks x of non-adopters are, by definition, unobservable. However, this additionally highlights the value of independent testing data in the context of app downloads, also highlighted more broadly, for example, by Romer, 2020, and Stock, 2020. For example, suppose the government were able to see adoption rates (e.g., tracing app downloads) by demographic group, as well as the results of existing infection tests by demographic group. A regression of adoption rates on infection rates at the group level would, given high enough quality of data, provide an approximation of

group-level adoption propensity scores.

Our framework can inform how best to target data collection efforts, through calculating adoption rate differentials by geography including local infection rates, across demographic groups and so forth. The model can help to highlight the location of “blind spots” in which information would confer public benefits, but is lacking because of inadequate adoption. This is important to achieve targeted and efficient public health efforts in such blind spots (e.g., information campaigns focused on adoption using social networks, GPs, local political networks, or other mechanisms).

4 Conclusions

From evidence available to date in the coronavirus pandemic, several countries appear to be struggling to achieve high adoption rates of government-endorsed contact tracing apps.⁴ We provide a simple theory framework to outline some trade-offs associated with the rollout of such tracing apps to acquire information to support policy responses. Our framework highlights the interplay between the representativeness of the data acquired from these apps, and the net benefits perceived by users in engaging in app downloads. The framework reveals that widespread adoption will depend on a number of factors including the security of the data collection exercise, and the possibility of anonymity guarantees for adopters, and helps to guide discussions on possible directions for policy. Overall, we believe this area is important and generally understudied, and hope that our initial contribution in the form of a framework will help to spark more involved efforts on these questions going forward.

⁴According to this blogpost, Covid-19 contact tracing apps have been installed by 9.3 percent of residents across the world’s 13 most populous countries with government-endorsed apps. In Italy, one of the worst hit countries in the first wave of the pandemic, the percentage, for instance, was 7.2%, while it was 3.1% in France (July 2020 figures).

References

1. Acquisti, A., L. Brandimarte, and G. Loewenstein (2015), “Privacy and human behavior in the age of information,” *Science*, 347, 509–514.
2. Atkeson, A. (2020), “What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios,” NBER Working Paper 26867.
3. Berger, D., K. Herkenhoff, and S. Mongey (2020), “An SEIR Infectious Disease Model with Testing and Conditional Quarantine,” Becker-Friedman Macro Finance Research Program Working Paper 2020-25.
4. Bordalo, P., N. Gennaioli, and A. Shleifer (2012), “Salience in Experimental Tests of the Endowment Effect,” *The American Economic Review*, 102 (3), 47–52.
5. Eichenbaum, M., S. Rebelo, and M. Trabant (2020), “The macroeconomics of epidemics,” NBER Working Paper 26882.
6. Goldfarb, A. and C. Tucker (2011), “Online display advertising: Targeting and obtrusiveness,” *Marketing Science*, 30, 389–404.
7. Goldfarb, A. and C. Tucker (2012), “Shifts in privacy concerns,” *American Economic Review*, 102, 349–53.
8. Morewedge, C. K., L. L. Shu, D. T. Gilbert, and T. D. Wilson (2009), “Bad riddance or good rubbish? Ownership and not loss aversion causes the endowment effect,” *Journal of Experimental Social Psychology*, 45 (4), 947–951.
9. Radaelli, L., Sapiezynski, P., Houssiau, F., Shmueli, E. and Y.A. de Montjoye (2018), Quantifying surveillance in the networked age: Node-based intrusions and group privacy. arXiv preprint arXiv:1803.09007
10. Romer, P. (2020), Roadmap to Responsibly Reopen America.
11. Stock, J. H. (2020), “Data Gaps and the Policy Response to the Novel Coronavirus,” NBER Working Paper 26902.
12. Varian, H. R. (2009), “Economic aspects of personal privacy,” in *Internet policy and economics*, Springer, 101–109.

13. Varian, H. R. (2010), "Computer mediated transactions," *American Economic Review*, 100, 1–10.
14. Vaudenay, S. (2020), Analysis of DP3T. *IACR Cryptol. ePrint Arch.*, 2020, p.399.