How To Make Digital Proximity Tracing Work: The View from Economics[1]

Jean-Philippe Bonardi[2], *Marius Brülhart*[3], *Jean-Pierre Danthine*[4], *Anshuman Saxena*[5], *Christian Thöni*[6], *Mathias Thoenig*[7], *Christian Zehnder*[8]

June 24th, 2020

*** *Preliminary version* ***

Abstract:

Digital proximity tracing (DPT) holds great promise as a complement to traditional proximity tracing, mainly because it allows contacts in anonymous crowds to be alerted in addition to known contacts. However, even the best technological solution, when it depends on widespread adoption, needs to contend with the human factor. Recent surveys in Switzerland suggest growing skepticism. We therefore assess the incentive structure faced by users of decentralized DPT apps such as SwissCovid. We show that DPT is impossible to enforce by decree or material incentives. Its success therefore hinges on widespread voluntary cooperation. Compliance essentially breaks down into 'passive' actions (downloading the app and carrying it around) and 'active' actions (entering an alert when tested positive). We show that 'passive' adoption generates considerable private benefits that many may underestimate – especially by offering information to guide their behavior towards vulnerable relatives and friends. 'Active' compliance is closer to a typical social dilemma, where private costs must be weighed against societal benefits. We argue that the costs are likely to be negligible for most users, while the social benefits (saving lives) are potentially large. Eliciting the necessary pro-social motivation among a majority of people will crucially determine the potency of DPT in the fight against the pandemic.

[1] We are grateful for helpful comments to Ed Bugnion and Blaise Genton.

[2] University of Lausanne (Department of Strategy, HEC Lausanne) and E4S; jean-philippe.bonardi@unil.ch
[3] University of Lausanne (Department of Economics, HEC Lausanne) and E4S; marius.brulhart@unil.ch
[4] EPFL (College of Management of Technology) and E4S; jean-pierre.danthine@epfl.ch
[5] EPFL, E4S and Synaestech; anshuman.saxena@synaestech.com
[6] University of Lausanne (FDCA) and E4S; christian.thoeni@unil.ch
[7] University of Lausanne (Department of Economics, HEC Lausanne) and E4S; mathias.thoenig@unil.ch
[8] University of Lausanne (Department of Organizational Behaviour, HEC Lausanne) and E4S; christian.zehnder@unil.ch

1. Introduction: The Promise of Digital Proximity Tracing

Contact tracing is an epidemiological intervention aimed at identifying people that come into close contact with a person suffering from a contagious disease. The widespread adoption of digital technologies, particularly mobile phones, for communication purposes, provides an opportunity to complement the labor-intensive, traditional methods with Digital Proximity Tracing (DPT). DPT allows contacts to be alerted immediately once an infection is detected, without the inevitable delays associated with personalized contact tracing through the healthcare system. A successful implementation of DPT would (i) help prevent the outbreak of a second wave of the Covid-19 pandemic (public health objective) and (ii) help avoid a costly reactivation of lockdown measures (economic objective).

DPT is particularly useful for identifying potential chains of infection in anonymous crowds such as in public transport, large events or nightlife venues, where standard contact tracing is ineffective. There are indications that the presence of 'superspreaders' and/or 'superspreading circumstances' at large gatherings could be a particularly important source of transmission in the case of Covid-19 (Fukui & Furukawa 2020). This would strengthen the case for complementing human contact tracing with a digital approach. DPT can moreover be useful in the case of a very large outbreak that overwhelms manual contact tracing organizations.

DPT is a paradigmatic example of a clever technological innovation that might ultimately fail because of mistrust and non-adoption by users (see Box 1 below). Although DPT can be considered useful as soon as one infection is successfully signaled, in order to deploy its full potential, DPT needs to be adopted by a large proportion of the population – 80% of smartphone users and 56% of the population, according to a recent study (Hinch et al. 2020). The success of DPT therefore not only requires a well-functioning technology (the engineering problem), but also broad acceptance among potential users (the social problem). Approaches from economics, and behavioral social science more generally, can help develop solutions to overcome the latter problem.

From an economist's point of view, choices are analyzed through the lens of *cost-benefit analysis*. The idea is that people take any action whose benefits exceed its costs – where these benefits and costs can be real or merely perceived.

In the case of DPT adoption decisions, cost-benefit analysis is complicated, because the costs are fully private, whereas the benefits are partly private and partly public. In this sense, DPT adoption is a version of the public goods problem (Samuelson 1954, Ostrom and Ostrom 1977). To be precise, we distinguish between three types of benefits that a DPT adoption decision of an individual creates: private benefits, public benefits to proximate others (family and friends), and public benefits to broader society.

The central question is therefore how to motivate people to bear the private costs of a decision for which a considerable part of the benefits goes to others.

This basic set-up, however, is further complicated by two additional considerations, with both positive and negative potential impacts on the adoption decision.

First, DPT is by its nature a class of product that economists associate with 'network effects'. This means that the value of the product is significantly driven not only by the product itself and the quality of the technology underneath, but also by the number of other individuals who have adopted the same product (hence our assertion that widespread adoption is required for full effectiveness). In the case of DPT, these network effects arise from the fact that the likelihood of a signal regarding a potential infection depends on whether a large share of the population has adopted the system in the first place. This implies that the benefits of DPT increase non-linearly with the number of users, and it strengthens the case for fast and widespread adoption of DPT. We return to this issue below.

Second, truly adopting the DPT technology is not a single choice but consists of a series of choices. In order for the technology to be fully effective, it is crucial that people not only install the app, but also that they actually carry their phone with them when they meet other people, that they respond adequately when they receive a warning from the app, and that they enter the code in the app after a positive serologic test result. This multiplicity of interconnected choices turns the situation into a so-called *multi-tasking problem* (Holmstrom and Milgrom 1991, Baker 1992). For this type of problem, it is very important to consider that behavioral stimulants that target one particular choice in the sequence of choices may have undesired effects on other choices in the sequence.

Box 1: Cognitive biases in time of crisis and the virtue of public communication

The pandemic has been characterized by a high level of uncertainty and anxiety. In this context of intense stress, cognitive biases tend to be amplified, irrational beliefs flourish and judgment is affected. Particularly problematic is confirmation bias, where people only give credit to versions of the facts that validate their own beliefs – a self-confirmation loop that is the basic fuel of all conspiracy theories. Examples abound, from the controversy about the curative role of hydroxychloroquine to the various theories about the origins of the virus. In the case of DGT, cognitive biases and misperception could blur the rational analysis of costs and benefits underlying the decision to adopt the app.

The intelligibility of public action is one of the major challenges of this health crisis. After strong initial popular support for containment policies, distrust and incomprehension of policy actions are now gaining momentum. The problem is that, if the logic of public action is not well understood, adherence to precautionary measures will slacken and the pandemic will flare up again.

A key aim of this paper, therefore, is to catalogue and assess costs and benefits of DPT as comprehensively and objectively as possible.

2. Benefits, Costs, and Four Specific Actions

At the time of writing, many potential adopters seem to perceive the private cost of adopting the technology as being very high, whereas the perceived private benefits seem to be perceived as rather low. According to opinion polls, the share of Swiss residents who would be willing to install the DPT app has fallen from 65% in April to 54% in June (SRF 2020). This perceptional imbalance may lead to a low adoption rate and ultimately imperil the chances of the whole DPT initiative in Switzerland.

There exist essentially two ways for policy makers to affect the cost-benefit evaluations of potential users.

Box 2: Technological and Privacy Risks of Decentralized DPT

DPT involves two categories of privacy risks: those associated with 'passive' use of the app and those associated with entering an infection alert.

So far, the only cybersecurity risk that has been identified for DPT apps such as SwissCovid is the replay of anonymized IDs broadcasted over Bluetooth beacons. An malicious attacker could conceivably record a large number of anonymized IDs of future-positive people, say, from around a Covid-19 testing facility, and then broadcast them in another location, thereby corrupting the system with fake contacts and potentially generating false positives. For the individual user, the risk from a replay attack implies the inconvenience of receiving a false positive alert.

Importantly, DPT apps do not use geolocation data, which reveal more about individual movements than Bluetooth-based contact data. It would be inconsistent to refuse to install a decentralized DPT app out of privacy concerns whilst running any app that uses location services (mapping app, health app, location-based games, etc.).

Privacy concerns associated with entering an infection alert are more straightforward and not related to concerns about technology, cybersecurity or general data protection issues. If I post an alert, my health status may be revealed to some of my recent contacts. There is no stigma in contracting the virus, and infection is likely difficult (and not judicious) to conceal from close contacts anyway. Yet, this might be a concern in some (rare?) situations. In particular, by activating the app alert, one might reveal recent contacts whom one prefers to keep confidential. With DPT, contact A is not informed that the same signal is received by contact B, but this link could possibly be inferred if the infected person only has a small group of recent contacts. Note, however, that this risk also exists and is even more obvious in the case of traditional contact tracing, since there is a human being collecting the information. Hence, this privacy risk should not be held against DPT. Note as well that, if perceived as acute, this risk could easily be avoided by not activating the app during confidential contacts. All in all, while situations could be imagined where sending an alert through the app could reveal unwanted information, such scenarios seem rather improbable. On the other side of the scale, it should be clear that the stakes are really high. If I have been tested positive, concealing this information even to strangers could leads to suffering and deaths. Given the high stakes (saving lives), it would amount to a serious abdication of civic responsibility to refuse to alert one's contacts through the app.

First, policy makers can act on potential users' information sets. There are good reasons to believe that many people have biased perceptions of the true costs and benefits that are involved in this decision. Conspiracy theories, technical misunderstandings and other misperceptions can create the impression of a very high private cost of DPT adoption (see also Sunstein and Vermeule 2009), whereas many of the private benefits and the benefits to proximate others do not seem to have reached the attention that they deserve. Attempts to debias such distorted evaluations are therefore one possible route to go (see e.g. Morewedge et al. 2015). We summarize the main privacy risks in Box 2.[9]

Second, one can directly target the actual costs and benefits of the decision to adopt. Such interventions can come in a variety of forms such as direct changes to the app itself, making access to certain activities contingent on having an active app installed, or reducing the cost of acting upon the information received through the app (free testing, support for voluntary/mandated quarantine).

Before we can discuss behavioral and economic implications of DPT, we need to define what use of DPT actually means. It is important not to equate app downloads with DPT adoption. In fact, four actions are jointly necessary for the DPT system to work:

Action 1: Download the app

Action 2: Carry the app-enabled smartphone during social contact situations, with Bluetooth switched on

Action 3: React to warning alert received from the app by self-isolating, getting tested, etc.

Action 4: Alert the system upon receipt of a positive Covid-19 test result

In the following, we distinguish between what we refer to as 'passive' and 'active' use of the app. Actions 1 and 2 are passive in so far as they require no interaction with the health system, whereas Actions 3 and 4 involve actions to be undertaken with respect to health-relevant alerts.

For a schematic overview of actions and associated costs and benefits, see Table 1 at the end of this paper. The various elements of the incentives structure are discussed in detail in the text below.

3. Encouraging Passive Use of the App: Three Nudges and an Incentive

3.1 Why Voluntary Compliance Is Essential

Given the high stakes in a pandemic, forced adoption via government decree might appear at first sight as an attractive option. Enforcing the adoption of the DPT app by law may be the most effective way to achieve maximum compliance with respect to Action 1 (installing the app). However, apart from legal concerns and problems in terms of enforcing mandatory use of the application, a number of other problems speak against such strong measures.

First, the four required Actions leave room for a large number of evasion strategies. It is well known from the literature on incentives in multi-tasking environments that strong incentives for one activity can backfire in terms of lower effort devoted to the other activities (see e.g. Baker et al. 1994 for examples). It is thus likely that people who are only weakly motivated to contribute to DPT, when forced to install the application, have even less reason to comply with the remaining actions.

[9] Our analysis considers Bluetooth-based decentralized privacy-preserving proximity tracing apps (DP3T) such as SwissCovid (see Troncoso et al. 2020). Other app designs might imply different privacy issues from those we consider in this document.

Similar concerns also hold for interventions that mandate the use of the DPT application for specific activities, such as using public transport. As the process of DPT requires a large number of people to comply at every step of the process, a system based on strong incentives will not be feasible in a free society. We will therefore turn our attention to less invasive measures, which focus on *voluntary compliance*.

In contrast to legal requirements or strong (material) incentives, *nudges* may be resorted to, as soft, minimally invasive tools to motivate people to install the application and contribute to the functioning of DPT (see Box 3).

One of the key characteristics of a nudge is that the intervention preserves the full freedom of choice, i.e., in contrast to legal requirements or strong economic incentives, individuals are still free to do what they want without adverse legal or financial consequences (see Thaler and Sunstein 2003). While other interventions may be more powerful, the feature of preserving freedom of choice is a key advantage of nudging.

Box 3: Nudges

The concept of nudges was proposed by Thaler and Sunstein (2008) as an alternative to traditional legal and economic measures to change individual behavior. Much of the original literature on nudges focused on situations in which individuals seem to fail to make optimal choices on their own (e.g. saving enough for retirement), but change their behavior in response to seemingly innocuous changes of the decision situation. Nudges tend to work best when individuals are facing a conflict between different goals. For example, the choice of healthy over unhealthy food often comes with short run costs but long run benefits. An individual aware of the trade-off might respond strongly to subtle cues, like presenting the healthy option more prominently on the menu.

3.2 Nudge 1: Framing

In public discussions of DPT, two topics have attracted most attention so far: (i) concerns about privacy, and (ii) the potential benefits of DPT for society at large. This dichotomy is problematic, because it emphasizes the social dilemma character of the decision to adopt DPT. In the public perception, the decision to install the application is then similar, for example, to fighting global warming: any individual contribution (e.g. not flying) comes at a substantial private cost but has no material private benefit. In such an environment, there is little hope that the problem will be solved by voluntary individual initiative alone.

In the case of DPT, however, the tradeoff privacy-for-public health is a misconception. First, it seems that the cost side is inflated by exaggerated concerns about privacy (see Box 2). Most smartphone users are willing to share detailed individual data with large international companies, compared to which any additional loss in privacy due to DPT should be negligible.[10] A nudging strategy should aim at reducing the perceived costs of Actions 1 and 2.

[10] The fact that DPT involves government in addition to private firms may heighten certain potential users' concerns. This ought to be taken into account both in designing the system and in communication. The fact that both powerful global firms and individual governments are involved might also present an opportunity, insofar as each may be acting as a check on the other.

→ Insight 1: Information campaigns in connection with the rollout of the DPT app should address privacy risks by assessing them relative to those of other common smartphone apps.

On the benefit side of the decision, it is important not to focus exclusively on the public-good character of DPT. While the literature in behavioral economics has shown that a large fraction of the population has prosocial inclinations, there remains no doubt that individual interests are the predominant motive for most people (Fehr and Schmidt 2003, Cooper and Kagel 2013). Public communication in connection with the DPT application should therefore highlight that there are private benefits to using the DPT application. Purely private benefits (to the individual him/herself) arise because knowing early about their exposure to the virus gives people the option of adapting their behavior. They may for instance choose to avoid stressful and tiring activities, such as not to weaken their antiviral resistance. Or they may choose to not to travel to places with inferior or more expensive healthcare services. In the future, as medical protocols and treatments evolve, early risk information could also potentially allow access to preventive treatments.

→ Insight 2: Information campaigns should make it clear that using the DPT application also has substantial private benefits.

For most people, the key potential benefit of DPT will likely arise from the information it can offer underlying the assessment of whether contact with vulnerable relatives is advisable or not. Information about the risk one poses to potentially vulnerable relatives and friends will likely be perceived as very valuable by most people.

Indeed, experimental studies show that pro-social behavior is dramatically increased by social proximity: we tend to care a lot more about people we know than about random strangers.

Leider et al. (2009), for example, conducted online field experiments in large real-world social networks in order to decompose prosocial giving into three components: (1) baseline altruism toward randomly selected strangers, (2) directed altruism that favors friends over random strangers, and (3) giving motivated by the prospect of future interaction. They show that directed altruism increases giving to friends by 52% relative to random strangers, whereas future interaction effects increase giving by an additional 24% when giving is socially efficient.

A simple slogan could therefore highlight the importance of DPT in protecting people who are important to the individual: "Protect your loved ones"!

→ Insight 3: Information campaigns should emphasize the value of DPT as a means to protect ones' family and friends.

3.3 Nudge 2: Conditional cooperation

While nudges to focus people's attention on the private benefits may be sufficient for some, others will remain unconvinced that installing the app makes sense from a purely individualist perspective or even considering their friends and families. For these people, the decision situation will remain a trade-off between their selfish impetus (avoid the cost) and their desire to contribute to the public good (e.g. derive satisfaction from participating in a collective activity to fight the pandemic).

There is a large literature on based on experiments of such social dilemmas. One of the well-established results is that a majority of the population is conditionally cooperative (see e.g. Gaechter 2007). This means that they are willing to contribute, as long as sufficiently many others contribute as well. In a meta-analysis of the empirical literature, Thöni and Volk (2018) find that typically some 60 percent of experimental subjects are conditionally cooperative, while 20 percent act purely selfishly. Importantly for DPT, only a small fraction (some 4 percent) is unconditionally cooperative, i.e., willing to contribute irrespective of what others do. Consequently, the success of DPT might hinge on the participation of conditionally cooperative individuals, who install the application only if they believe that sufficiently many others use the application as well. Targeting information about the adoption rate towards conditionally cooperative individuals would then be helpful if it increases their expectation about the adoption rate.

The complicated part is expectations management. In the absence of any information about the adoption rate (e.g. prior to the rollout), people will form expectations about adoption by others. Informing the public about the fact that a certain percentage of the population installed the application might then come as a positive surprise to some, and motivate them to sign up as well. For others, the adoption rate might seem surprisingly low, and thereby reduce their motivation to sign up.

One option might be for the authorities to communicate along with daily new infection rates (i) the share of the tested cases that had received an app warning signal, and (ii) the number of these cases who were able to be communicated to recent contacts quickly through the DPT app. When presented in this context, high or growing shares of new infections detected by and/or notified via DPT would be a signal of success, and low or falling shares would be hard to frame (by the media) as anything else but an unsatisfactory trend, implying a kind of collective public shaming.[11] This might nudge the public discourse in the direction of encouraging higher adoption rates as well as encouraging active use of the app in case of infection (Action 4).

→ <u>Insight 4:</u> Adoption of the DPT app in the population could be communicated in terms of the daily share of new infections (i) that were detected thanks to an app signal received, and/or (ii) that led to an app signal sent by the infected person.

The concept of conditional cooperation is closely related to that of *network effects*. The standard framework for quantifying the relationship between private benefits and the size of the network Is Metcalfe's Law: the value of a network with n users is proportional to n^2. In other words, if a network is worth \$$x$ to a user for each user on the platform, 10 users on the platform would make the network worth \$100x to a user (Farrell & Saloner, 1985; Economides, 1996). The value of a network of DPT users should thus grow exponentially as the number of users in the network grows. Again, it is important that users realize and understand the growing benefits that are created by a growing number of DPT adopters. The millionth DPT adopter will add much more than one-millionth to the social value of the DPT system (assuming he/she does not have a significantly below-average frequency of social contacts).

→ <u>Insight 5:</u> Information campaigns could stress the fact that the social benefits of app usage rise exponentially with the adoption rate.

Human networks, however, have very irregular patterns. Within networks, nodal agents might inform and convince many others to adopt the app. Such 'influencers' could be targeted in order start growing the installed base of DPT users. Celebrities and social media influencers volunteering to contribute to the information campaign could significantly help propagation of app usage.

[11] This is comparable to recycling, where information about overall participation is sometimes used in a quest to strengthen a favorable social norm (e.g. Abbott et al. 2013).

Musicians, athletes and other celebrities whose professional activity directly depends on big public events (stadiums, concert halls, etc.) would be particularly credible advocates, highlighting that DPT is particularly useful for tracing cases in the public space with a high density of anonymous encounters.

3.4 Nudge 3: Opt-out instead of opt-in

One of the strongest nudges identified in the literature is people's tendency to adopt the default option. For example, the fraction of the population willing to donate organs can be increased when switching from an opt-in system to an opt-out (presumed consent) system, in which people have to actively state that they are not willing to donate their organs (Johnson and Goldstein 2003). While everybody is free to choose on organ donation, changing the default option increases effective donation rates dramatically.

A potentially powerful measure to foster the installation of the DPT app (Action 1) would be to link it to the normal updating process of the smartphone operating system. The system should still preserve the freedom to choose, by unchecking a box in case one does not want to install the app.

A somewhat less strong, but potentially similarly effective tool would be to implement an active choice environment. From the literature on organ donations, we know that the main reason for a low rate of donors is that people are unwilling to even think about the question (e.g. Beshears et al. 2020). If it were possible to create a situation in which users are forced to make a choice, then a much larger fraction would comply than in an opt-in scenario. For DPT this would mean that the normal updating process of the smartphone could ask the user whether or not they would like to install the application, thereby forcing a decision.

While making the installation of the app the default option would likely be the strongest nudging intervention, this option is not available at the moment, as Apple and Google have decided against it. However, if the pandemic were to flare up again and cause even greater harm than currently anticipated, this option may have to be reconsidered.[12]

→ Insight 6: Forcing users to deliberately opt out of installing the DPT app might be an option worth considering if the pandemic were to cause greater future damage than currently expected.

3.5 Incentive: free testing

As long as PCR testing is subject to supply constraints, the right to free testing upon receipt of an app warning could provide a powerful incentive to install and use the app. The prospect of free testing, as foreseen in the Swiss Covid-19 legal framework, both offers a material private reward and signals to app users that their action is socially valuable. Subsidies for testing are a very cost-effective policy measure in a pandemic (Bütler et al. 2020).

→ Insight 7: Testing upon receipt of an app warning should be free.

---

[12] Technically, the challenge would be to adapt operating systems in such a way that they generate an automatic opt-in by country of current residence. The active-choice option might therefore be more realistic, where users could choose from a list of official DPT apps. In Switzerland, compatibility with the legal provision that the official DPT app be "voluntary" would moreover need to be ascertained.

One might even consider paying alerted app users to get tested. While one might offer people a small amount designed to cover their travel expenses, more generous payment risks being counterproductive as it could crowd out intrinsic motivation (e.g. Gneezy et al, 2011), and it could also be interpreted as implying that getting tested is not privately advantageous for the concerned persons.

4. Encouraging Active Use of the App: a Delicate Balance

Assuming a successful launch of DPT with a sufficient percentage of the population installing and activating the app (Actions 1 and 2), the efficacy of DPT relies on users' reaction to the information received (Actions 3 and 4).

4.1 Push Hard Early or Wait and See?

Network externalities and adherence to a cultural norm are two social forces that trigger self-reinforcing adoption behaviors. If many in society at large use the app, then both the marginal private benefit of using the app and the reputational cost of not using it will be high; consequently, even more people will have an incentive to use the app. But this virtuous logic can revert and a vicious cycle of non-adoption can also be sustained in the long-run. Hence, the adoption rate can follow a multiplicity of trajectories and nothing guarantees that the focal point of behaviors will be on the high adoption trajectory. There may be considerable path-dependence in these trajectories. In particular, some early-stage accidental events could determine whether the society will engage into a low- or high- adoption trajectory (snowballing effect).

These observations have three implications:

- Because of snowballing, being pro-active in the beginning of the app rollout through a large information campaign could be crucial (coordination of expectations and behaviors).
- Conversely a "laisser-faire" approach at the introduction stage carries risks, because snowballing is associated with strong lock-in effects. According to that logic, expecting that the threat of a second wave in Fall 2020 will trigger a wave of adoption could be overly optimistic.
- With a multiplicity of trajectories, cross-country comparison of adoption rates can be misleading. A good adoption trajectory, say in Germany, is not a guarantee of success in Switzerland.

An alternative view is that network externalities may have less weight, and individual cost-benefit calculations matter more. If so, the heightened risks in a second wave may be sufficient to trigger a wave of DPT adoption even if the initial take-up had been low. This might even allow technological teething problems during the rollout to be corrected without causing too much disruption (given that early adopters likely are also more technology affine).

In the face of such uncertainties, it is unclear whether an early 'big push' public information campaign is to be preferred to a more low-key introduction.

4.2 *Beware of the "Cry Wolf" Problem*

The sensitivity of the DPT alert system is the crucial technical parameter in this problem. Like any medical indicator, there will be two kinds of errors: the error of false positives (the app indicates close contact with an infected individual, but no transmission occurred), and false negatives (the app is used correctly and the user gets infected, but no alert is sent).[13]

As long as the population adoption rate is below 100 percent, users will understand that false negatives cannot be held against the DPT itself. However, the same is not necessarily true for false positives.

From the medical point of view, false positives are less problematic than false negatives – it is less costly to isolate a person needlessly than not to isolate an infectious person. Even standard contact tracing is known to yield a preponderance of false positives, with actual infections among the alerted persons typically in the single-digit percentage range (Salathé et al. 2020).

While everyone understands that close and prolonged proximity with an infected person does not guarantee being infected, there is a substantial "cry wolf" problem if the system parameters are set too sensitively: repeated false positives will likely reduce compliance with Action 3 (see also Normile 2005). Consequently, the parameters on the sensitivity of the DPT alert system have to strike the right balance between false negatives and false positives. This implies setting a somewhat lower sensitivity than what would be advisable on purely medical grounds.

→ <u>Insight 8:</u> The parameters of the DPT app should not be set too sensitively, so as to mitigate the "cry wolf" problem.

This problem, however, is only likely to be relevant in times and places of very high virus prevalence.

4.3 *Facilitate Self-Isolation*

Upon receiving a warning from the app, conditional on having applied for a test, and until having received a test result, people should self-isolate. Even if the app were to produce a large number of false positives (in the sense that a proximate contact did occur, but the virus did not transmit), the precautionary principle suggests that people should self-isolate if possible. The probability of adherence to the self-isolation advice will be maximized if people need not fear any loss of income or discrimination in the workplace (Bütler et al. 2020). It is thus important that income-replacement payments be available without administrative complications to anybody who self-isolates between receiving an app warning and receiving a PCT test result. If they test positive, standard quarantine regulations will apply. If they test negative, the right to paid self-isolation ends.

→ <u>Insight 9:</u> Income-replacement payments and legal protections should be offered to people who self-isolate in the period between receiving an app warning and receiving their PCT test result.

4.4 *Action 4: An Almost Pure Social Dilemma*

Finally, there is Action 4: after having received a positive PCT test result, it is up to the app user to enter the code received from the medical personal into the system. Compliance at this final step

[13] False negatives could result not only from insufficiently sensitive technical settings but also from a failure by infected persons to enter the alert (Action 4).

E4S
Enterprise for Society

will be crucial. If a person who has tested positively and been in close contact with unknown others fails to enter the information into the app, this may result in avoidable additional infections. The public benefits are therefore potentially large. The personal benefits and the benefits to proximate others, in contrast, appear comparatively small. Upon testing positively, an app user has no purely selfish advantage from entering the information, and proximate others (such as friends and family) can be alerted directly through channels outside the app. A slogan such as 'protect your loved ones' therefore loses its appeal at this stage.

Hence, Action 4 presents an almost pure social dilemma, with potential private costs (see Box 2) and predominantly social benefits.

There are two insights from the public goods literature that are of importance here. First, people are more likely to contribute in social dilemma situations if the private cost of contributing is small relative to the public benefit (Isaac and Walker 1988). The direct cost of entering a positive test result in the app is negligible, but people might be concerned about being held responsible for disseminating the virus if alerted persons can identify where the signal came from. Such an identification might be possible if some of the alerted persons had only recent contact to a very small group of people. This would only be a real concern to the extent that the infected person did not want some of their contacts to be revealed to each other. Moreover, should this be the case, not entering the information may well not protect the decision maker. If the other person is indeed infected and shows symptoms, the disseminator could also be identified without the app.

In contrast to those costs, the public benefit is potentially very large: it is literally a matter of life and death.

→ Insight 10: Public messaging about app usage should emphasize the 'life-and-death' nature of the public benefits of entering information, relative to privacy concerns that are unlikely to be important for the vast majority of people.

Second, peer punishment is a powerful mechanism to sustain contributions to public goods (Fehr and Gächter 2000). Many people punish non-contributors even if doing so is costly. Punishment is effective in enhancing altruism, even when it comes in form of verbal criticism rather than a material cost (Bowles and Gintis 2002, Ellingsen and Johannesson 2008). It seems that (verbal) sanctions evoke emotions of shame and induce people to comply. The problem with using versions of punishment in the context of Action 4 is that non-contributors cannot be personally identified. However, it would be possible to observe the overall contribution rate, the number of codes issued and the number of alerts entered into the system were collected. Such information could then be used to encourage people to comply with Action 4.

→ Insight 11: Compliance rates for information entry could be published, combined with messaging about the potentially large cost of non-compliance. Moreover, medical units that perform PCT tests should strongly advise users to enter their codes after a positive result, e.g. by following up with a phone call.

If non-compliance with Action 4 turned out to be a big problem, one might also think of turning alert entries from an opt-in choice, as currently foreseen in Switzerland, to an opt-out choice. Making Action 4 mandatory, however, would risk undercutting compliance with Actions 1-3.

5. Conclusions

DPT is generating much excitement at a time when the world is crying out for ways to contain the Covid-19 virus. It is a clever but barely tested technological approach that could offer a useful complement to standard contact tracing. DPT may end up delivering less than its most optimistic proponents are expecting, but it is a cheap measure that is highly unlikely to cause any harm and has upside potential for saving lives and avoiding unnecessarily crude lockdowns.

However, even the most brilliantly engineered technological solution, when it depends on widespread adoption, needs to contend with the human factor. In this paper, we provide a first attempt at doing that.

We have considered the individual parts of users' decision problem, described the associated incentive structures, and highlighted some policy-relevant insights. We argue that users need to adopt a number of behaviors for DPT to be effective, and that these behaviors can neither be decreed nor effectively enforced in a free society. The success of DPT therefore hinges on the elicitation of widespread voluntary cooperation.

The main elements of our analysis are summarized in Table 1. We catalogue costs and benefits associated with each of the four user actions that are necessary for successful adoption of the DPT technology.

It turns out that in terms of incentives for adoption, the most meaningful distinction is between Actions 1 to 3, which do not require entering any information into the app, and Action 4, which consists of entering an alert upon receiving a positive test result.

The first three actions, looked at objectively, do not pose much of a dilemma. Very small costs in terms of privacy risks stand against substantial private as well as social benefits. Probably the main private benefit is the information the app can offer to guide users' contact decisions with respect to vulnerable family members and friends: information about one's likelihood of being contagious will be valuable primarily for guiding behavior towards close others. Given people's much documented propensity to act much more 'pro-socially' towards relatives and personal acquaintances than towards strangers, it could be important in official communication to make people aware that using the app brings advantages to themselves in addition to contributing to the public good. As a complement to pointing out such private benefits, publication of app-based notification rates could moreover nudge society towards establishing a norm whereby not using the app is considered an anti-social act.

The fourth action, notifying an infection, poses a more difficult incentive problem. Some people may have entirely rational privacy concerns about communicating one's health status to all their recent contacts, and the private benefits appear to be quite limited (alerting some close contacts one might otherwise have forgotten about). However, the benefits to strangers are potentially very large: one might save the lives of people one sat next to in a crowded space – or of their vulnerable contacts. This is therefore close to a pure 'social dilemma', where people have to pay a private cost to provide an advantage to faceless strangers.

While DPT adoption may therefore be promoted to some extent by appealing to people's selfish instincts, full compliance will not be achievable without some degree of public-spiritedness on the part of all. Ultimately, the success of DPT will rely on the better angels of our nature.

References

Abbott, A., Nandeibam, S., & O'Shea, L. (2013) Recycling: Social Norms and Warm-Glow Revisited. *Ecological Economics*, 90, 10-18.

Baker, G. P. (1992). Incentive Contracts and Performance Measurement. *Journal of Political Economy*, 100(3), 598-614.

Baker, G., Gibbons, R., & Murphy, K. J. (1994). Subjective Performance Measures in Optimal Incentive Contracts. *Quarterly Journal of Economics*, 109(4), 1125-1156.

Beshears, J., Choi, J.J., Laibson, D., & Madrian, B.C. (2020). Active Choice, Implicit Defaults, and the Incentive to Choose. *Organizational Behavior and Human Decision Processes*, forthcoming.

Bowles, S., & Gintis, H. (2002). Homo Reciprocans. *Nature*, 415(6868), 125-127.

Bütler, M. et al. (2020) Economic Considerations of Test-Isolate-Trace-Quarantine (TiTQ). *Policy Brief*, Swiss National COVID-19 Science Task Force.

Cooper, D., & Kagel, J. (2013). Other-Regarding Preferences: A Selective Survey of Experimental Results. In J. H Kagel and A. E Roth (eds), *Handbook of Experimental Economics*, Vol 2, Princeton University Press.

Economides, N. (1996). The Economics of Networks. *International Journal of Industrial Organization*, 14, 673-699.

Ellingsen, T., & Johannesson, M. (2008). Anticipated Verbal Feedback Induces Altruistic Behavior. *Evolution and Human Behavior*, 29(2), 100-105.

Farrell, J., & Saloner, G. (1985). Standardization, Compatibility, and Innovation. *Rand Journal of Economics*, 16(1), 70-83.

Fehr, E., & Gächter, S. (2000). Cooperation and Punishment in Public Goods Experiments. *American Economic Review*, 90(4), 980-994.

Fehr, E., & Schmidt, K. (2003). Theories of Fairness and Reciprocity: Evidence and Economic Applications. In M. Dewatripont, L. Hansen, & S. Turnovsky (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, Eighth World Congress (Econometric Society Monographs, pp. 208-257). Cambridge: Cambridge University Press.

Fukui, M., & Furukawa, C. (2020). Power Laws in Superspreading Events: Evidence from Coronavirus Outbreaks and Implications for SIR Models. *Mimeo*, Department of Economics, Massachusetts Institute of Technology.

Gächter, S. (2007). Conditional Cooperation: Behavioral Regularities from the Lab and the Field and their Policy Implications. In B. S. Frey & A. Stutzer (Eds.), *Economics and Psychology: A Promising New Cross-Disciplinary Field* (p. 19–50). MIT Press.

Gneezy, U., Meier, S., & Rey-Biel, P. (2011). When and Why Incentives (Don't) Work to Modify Behavior. *Journal of Economic Perspectives*, 25(4), 191-210.

Hinch, R., et al. (2020). Effective Configurations of a Digital Contact Tracing App: A Report to NHSX. *Policy Report*, University of Oxford.

Holmstrom, B., & Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law Economics & Organization*, 7, 24.

Isaac, R. M., & Walker, J. M. (1988). Group Size Effects in Public Goods Provision: The Voluntary Contributions Mechanism. *Quarterly Journal of Economics*, 103(1), 179-199.

Johnson, E. J., & Goldstein, D. G. (2003). Do Defaults Save Lives? *Science*, 302, 1338–1339.

Leider, S., Möbius, M.M., Rosenblat, T., & Do, Q. (2009). Directed Altruism and Enforced Reciprocity in Social Networks, *Quarterly Journal of Economics*, 124(4), 1815–1851.

Morewedge, C. K., Yoon, H., Scopelliti, I., Symborski, C. W., Korris, J. H., & Kassam, K. S. (2015). Debiasing Decisions: Improved Decision Making with a Single Training Intervention. *Policy Insights from the Behavioral and Brain Sciences*, 2(1), 129-140.

Normile, D. (2005). Pandemic Skeptics Warn Against Crying Wolf. Science, 310(5751), 1112-1113.

Ostrom, V., & Ostrom E. (1977), Public Goods and Public Choices, in E. S. Savas (ed.), *Alternatives for Delivering Public Services: Toward Improved Performance*, Boulder, CO: Westview Press, 7-49.

Salathé, M., von Wyl, V., Vayena, E., & Bugnion, E. (2020). Digital Proximity Tracing. *Policy Brief*, Swiss National COVID-19 Science Task Force.

Samuelson, P. (1954), The Pure Theory of Public Expenditure, *Review of Economics and Statistics*, 36, 387–389.

SRF (2020). *Die Angst vor dem Coronavirus verfliegt: 4. Corona-Umfrage der SRG*. www.srf.ch/news/schweiz/4-corona-umfrage-der-srg-die-angst-vor-dem-coronavirus-verfliegt

Sunstein, C. R., & Vermeule, A. (2009). Conspiracy theories: Causes and cures. *Journal of Political Philosophy*, 17(2), 202-227.

Thaler, R. H., & Sunstein, C. R. (2003). Libertarian Paternalism. *American Economic Review*, 93(2), 175-179.

Thaler, R.H. & Sunstein, C.R. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. (Yale University Press).

Thöni, C., & Volk, S. (2018). Conditional Cooperation: Review and Refinement. *Economics Letters*, 171, 37-40.

Troncoso, et al. (2020). Decentralized Privacy-Preserving Proximity Tracing. *Mimeo*, arXiv:2005.12273.

Table 1: Overview of Actions and Associated Costs and Benefits from the Point of View of an Individual User of a DPT App

Actions	Costs	Benefits		
	to individual	*to individual*	*to family and friends*	*to society at large*
1. Download app	None	Better information on infection risk ⇒ • adapt own plans (e.g. travel) • adapt own health behavior (e.g. stress levels)	Better information on infection risk ⇒ adapt contacts with vulnerable 'loved ones'	• Contribution to containment of pandemic • Contribution to avoidance of renewed general lockdown • Exponential social value of individual participation due to network effect
2. Carry app around	• Battery drain • Risks associated with use of Bluetooth • Exposure to replay attack (false positive, sent maliciously)			
3. React to warning signal	• Material and psychological cost of distancing and self-isolation • Ex post psychological cost of overreaction in case of false alerts ('cry wolf')	• Free test • Reduced risk of social stigma as a spreader	Protect 'loved ones' by adapting contacts	
4. Enter warning signal when infected	Potential revelation of contact patterns and health status (privacy risk)	None	Alert and protect 'loved ones' quickly, easily and comprehensively (no imperfect recall problem)	Alert and protect unknown contacts