

Research



Cite this article: Barrat A, Cattuto C, Kivelä M, Lehmann S, Saramäki J. 2021 Effect of manual and digital contact tracing on COVID-19 outbreaks: a study on empirical contact data. *J. R. Soc. Interface* **18**: 20201000. <https://doi.org/10.1098/rsif.2020.1000>

Received: 9 December 2020

Accepted: 13 April 2021

Subject Category:

Life Sciences—Physics interface

Subject Areas:

biocomplexity, biomathematics, biophysics

Keywords:

COVID-19, contact tracing, temporal contact networks

Author for correspondence:

A. Barrat

e-mail: alain.barrat@cpt.univ-mrs.fr

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.5401758>.

Effect of manual and digital contact tracing on COVID-19 outbreaks: a study on empirical contact data

A. Barrat^{1,2}, C. Cattuto^{3,4}, M. Kivelä⁵, S. Lehmann⁶ and J. Saramäki⁵

¹Aix Marseille Univ., CNRS, CPT, Turing Center for Living Systems, Université de Toulon, Marseille, France

²Tokyo Tech World Research Hub Initiative (WRHI), Tokyo Institute of Technology, Tokyo, Japan

³Computer Science Department, University of Turin, Turin, Italy

⁴ISI Foundation, Turin, Italy

⁵Department of Computer Science, Aalto University, Aalto, Finland

⁶Technical University of Denmark, Copenhagen, Denmark

AB, 0000-0001-8683-269X

Non-pharmaceutical interventions are crucial to mitigate the COVID-19 pandemic and contain re-emergence phenomena. Targeted measures such as case isolation and contact tracing can alleviate the societal cost of lockdowns by containing the spread where and when it occurs. To assess the relative and combined impact of manual contact tracing (MCT) and digital (app-based) contact tracing, we feed a compartmental model for COVID-19 with high-resolution datasets describing contacts between individuals in several contexts. We show that the benefit (epidemic size reduction) is generically linear in the fraction of contacts recalled during MCT and quadratic in the app adoption, with no threshold effect. The cost (number of quarantines) versus benefit curve has a characteristic parabolic shape, independent of the type of tracing, with a potentially high benefit and low cost if app adoption and MCT efficiency are high enough. Benefits are higher and the cost lower if the epidemic reproductive number is lower, showing the importance of combining tracing with additional mitigation measures. The observed phenomenology is qualitatively robust across datasets and parameters. We moreover obtain analytically similar results on simplified models.

1. Introduction

The coronavirus disease (COVID-19) epidemic was declared a pandemic by the World Health Organization on 11 March 2020. As of mid-March 2021, COVID-19 has reached at least 120 million confirmed cases worldwide and caused at least 2.6 million deaths [1].

In the first half of 2020, faced with an exponentially growing number of cases and given the absence of effective pharmaceutical treatment and lack of vaccine, governments have first had to rely on broad, nationwide measures to reduce the mobility and number of contacts between individuals, starting with school closures and eventually lockdowns of whole countries [2–5]. These non-pharmaceutical interventions (NPIs) have succeeded in limiting contagion [6] and they have been gradually lifted in many cases after the first wave. However, the building of population immunity has been slow [7], and new epidemic waves have occurred. Overall, the situation appears under control in only a few countries, and many are experiencing successive ‘waves’ and re-emerging outbreaks, with a corresponding heavy burden on the healthcare systems and renewed regional or nationwide lockdowns. Exit strategies after the end of a lockdown are thus clearly needed to avoid new resurgences until vaccination has reached sufficient fractions of the population.

General measures to reduce transmission risk have remained in place or have been lifted and then re-imposed. These include enhanced hand hygiene, mandatory mask wearing in a number of contexts like public transportation,

shops or even whole cities and limits to the size of social gatherings. These measures are beneficial and very useful instruments and have in part mitigated the successive waves of the COVID-19 pandemic [8]. They are also much easier to enforce than regional or country-wide lockdowns and they have smaller economic costs. However, these measures have proven to be insufficient on their own. They are also still indiscriminate, as they are imposed on infected and non-infected individuals equally. Targeted measures, on the other hand, focus on stopping the spread where and when it occurs. The most common targeted intervention consists in the isolation of infectious individuals, which prevents the infected individual from further spreading the disease. This can be fully effective only if all infectious individuals are identified *before* they infect others (e.g. through symptoms or testing). However, in the case of COVID-19 a large fraction of transmissions occur before symptom onset [9–13]. Furthermore, many infected people develop only mild symptoms or no symptoms and are therefore typically not tested, despite being able to transmit the disease [12,14]. A natural way to cast a wider net around an identified case is the so-called *contact tracing* process, which endeavours to identify people who have been in long enough contact with an infectious individual [15,16]. The rationale is that those individuals, even in the absence of symptoms, might be infectious as well and that quarantining them and monitoring their health can both protect them and block further transmission.

Contact tracing is traditionally carried out via interviews of identified cases, followed by phone calls to the identified contacts to warn them and to ask them to go into quarantine [15,16]. Such ‘manual’ contact tracing (MCT) is labour intensive; it can be slow, and it critically relies on the ability of individuals to remember and identify their contacts. While this is easy in some contexts, such as households, our actual ability to recall and identify close-range proximity contacts is known to be limited. Retrospective surveys have found that brief contacts have a lower probability of being recalled and that contact durations are overestimated [17]. Moreover, in important contexts such as public transportation, shops or lifts (elevators), we often find ourselves in the proximity of unknown persons. In such cases, digital proxies for close-range proximity are currently viewed as a complementary and scalable approach, known as *digital contact tracing* (DCT), that could overcome the above limitations [18–20].

The broad adoption of smartphones has provided an opportunity to develop apps that use short-range device-to-device communication between phones to sense the proximity of their owners. The exchange of low-power Bluetooth packets between smartphones can indeed be used to detect proximity relations between persons whose phones are running the app: whenever a person who has installed a contact tracing app is diagnosed as infectious, a warning can be sent to all app users that the person has been in close proximity with during the previous few days [18,21,22]. Note that the proximity relations detected by Bluetooth are approximate and partially depend on the environment and detection thresholds that need to be set after calibration [21,23]. The wide dissemination of Bluetooth, however, currently makes it the best available proxy for physical proximity in DCT, and different technical architectures for DCT have been proposed, including so-called ‘decentralized’ solutions [21] that by design minimize the amount of information that is collected and shared. Exposed individuals warned via DCT are asked to contact the responsible health

authorities and to quarantine themselves. Importantly, contact tracing apps do not rely on individuals recalling or naming their contacts nor do they require knowledge of their identity. However, the efficiency of contact tracing apps is obviously limited by their numbers of users, since both the infectious person and their contacts need to have installed the app, and by the compliance of alerted users. Some studies [22,24] have proposed to further extend the reach of DCT by using information on individuals situated two hops away from a confirmed case, along the digitally sensed proximity network (so-called ‘recursive’ contact tracing). This might however have privacy implications, and it is important to note that the success of app-based contact tracing depends on a complex interplay of proximity tracing technology, citizens’ trust and adoption [25], effectiveness of the exposure notification strategy and good integration with traditional contact tracing and with other public health capabilities and processes.

Several theoretical studies have investigated the potential effectiveness of proximity tracing apps in the fight against COVID-19 [15,16,18–20,22,26–30]. The studies that have most influenced the public discussion around contact tracing apps are based on a simplified, macro-level mathematical model that assumes homogeneous mixing of the population [18,31]. However, real social and contact networks are highly non-homogeneous and exhibit many non-trivial structural and temporal features, such as heterogeneous distributions of contact durations [17,32,33], sets of ties with correlated activities [34] and intermittent communities and dynamical social structures on multiple time scales [35,36]. Heterogeneities in contact structure and durations are known to be highly relevant for epidemic spreading processes, as well as group structures [37,38]. In particular, they have been shown to influence the estimation of the epidemic size for simplified models of infectious disease dynamics [39,40]. However, it is at this stage unclear whether the estimation of an intervention’s efficiency is impacted by these complex structural and temporal properties of real contact patterns: in particular, the dependence of an intervention’s effect on its parameters (here, for instance, the efficiency and speed of contact tracing or the adoption rate of the digital tracing app) might *a priori* differ in simple contact models and in real contact networks. Therefore, to go beyond homogeneous mixing results, some studies have considered artificial static network structures to study at a theoretical level the effect of heterogeneous structures [24], and other approaches have used large-scale agent-based models [27,30] or GPS data [20] to recreate synthetic daily aggregated networks of contacts and to evaluate the effectiveness of app-based contact tracing in more realistic settings.

Using high-quality empirical data on close-range proximity between individuals is thus an additional crucial step to deepen our theoretical understanding of contact tracing, understand the phenomenology of contact tracing efficiency in real-world settings and thus improve the grounding of the discussion on the relative and combined efficiency of MCT and DCT. State-of-the-art datasets of this kind were collected and made public over the years by several research efforts, in particular by two independent collaborations [33,41] who have used proximity sensors and smartphones in specific communities (such as among university students) and in highly relevant social contexts such as workplaces, schools and hospitals. These empirical data are temporally resolved and exhibit the complex structural and temporal complex features mentioned above [32,33,35,36,42]. They represent thus an important

benchmark for realistic simulations of epidemic processes and interventions. We note here that these datasets were collected in non-pandemic situations. No such data are currently available to investigate how contact patterns have changed owing to restrictions and individual risk perception [43], and this has thus typically been modelled by effectively reduced transmission rates [4,29,30]. We consider here post-lockdown scenarios in which contact patterns and, behaviours could return to (almost) normal, and for completeness, we will also consider both reduced effective transmission rates as well as other ways to take such potential changes into account, such as random reduction of contacts or suppression of fleeting contacts. We also note that these data have been recently used [29] to derive parameters for the approach described in [18], and to determine—within that approach—which app-based contact tracing policies might be most effective.

Here, we use such empirical high-resolution contact network data to inform a compartmental model for COVID-19 [4] and to simulate directly the targeted measures described above in realistic, albeit circumscribed, contexts. In particular, we study the combination of app-based contact tracing and MCT, taking into account specific limitations of both approaches. We quantify the impact of interventions by measuring the reduction of the final size of the simulated epidemic. Moreover, we study the number of quarantines and the fraction of quarantined individuals who were not infected. We find that the mere isolation of symptomatic cases is not sufficient to substantially reduce the epidemic size, while interventions guided by contact tracing can have a strong impact. In particular, MCT, even when imperfect (i.e. limited recall of past contacts, delays in alerting contacts), yields a potentially strong reduction in epidemic size that grows linearly with the fraction of contacts correctly recalled. The effect of DCT, in turn, grows only quadratically with the fraction of app adopters, as expected from the constraint that both the case and the contact need to be running the app for a contact to be detected by their phones. Therefore, if DCT is used in isolation, a large fraction of adopters is required to obtain a large impact, although no threshold effects are observed: no specific adoption threshold exists that would lead the app to have no impact at a low level of adoption but suddenly become efficient if a certain level of adoption is reached. On the contrary, any degree of adoption yields a positive contribution. While the specific values of the epidemic size reduction due to contact tracing efforts depend on the parameters and contexts, the overall functional shapes are similar in all cases. In fact, we recover the overall qualitative behaviour of the epidemic size reduction in a simplified analytical model of propagation, showing the generality of this phenomenology.

We also observe that MCT and app-based contact tracing lead to similar numbers of quarantined individuals (which can be interpreted as the cost of the intervention) and fractions of non-infected quarantined individuals, for a given relative reduction in epidemic size (benefit of the intervention). In addition, even if the specific values obtained depend on parameters and contexts, this cost–benefit curve has a robust parabolic shape and exhibits a maximum: if contact tracing is effective enough, it might contain an outbreak early on, leading to few quarantine events, i.e. high benefit with low cost. We also provide analytical arguments to understand theoretically this robust parabolic shape.

Finally, we show that, for fixed app adoption and MCT parameters, recursive contact tracing further reduces epidemic

size, at the cost however of a higher fraction of non-infected individuals among quarantined individuals. In fact, even at given benefit (epidemic size reduction), the cost in terms of number of quarantines and non-infected quarantined individuals is higher for recursive tracing than for standard tracing. While the data we use are limited to specific contexts, the robustness of the observed qualitative behaviour in various contexts and with respect to parameter changes, together with the fact that we can understand them using analytical approaches, hints at the generality of our results. We conclude that DCT and MCT complement each other, and their combination can help reach an optimal regime of high efficiency where the spread is suppressed at low cost in terms of quarantines. Importantly, DCT compensates for the inherent limitations in recall and scalability of MCT, but improvements in MCT efficiency can yield particularly strong effects and should thus be prioritized. Moreover, app adoption should clearly be as large as possible, but there is no specific level of adoption that must be achieved, as any improvement in adoption has a positive impact.

2. Results

We model an outbreak of COVID-19 by means of a compartmental model [4] in which individuals can be in a series of discrete states describing the unfolding of the disease (see figure 1*a* and Methods for details and parameter values). Susceptible (healthy) individuals (S), upon contact with infectious ones, can contract the disease and first enter the exposed (non-infectious) state (E) and then a pre-symptomatic infectious state (I_p). Pre-symptomatic individuals can remain asymptomatic during the whole infectious phase (I_a), with probability a , or they can develop either mild or severe symptoms (I_m or I_s), with respective probabilities m and $s = 1 - a - m$. The infectious state leads to recovery R (or death) after a typical time.

We simulate the model on empirical data describing time-resolved, close-range proximity interactions of individuals in two relevant settings: a community of students (SD , in the following) and an office (OD , in the following). The SD dataset describes the network of physical proximity among a population of more than 700 university students over a period of four weeks. The dataset was collected as part of the larger Copenhagen Networks Study [23,33], in which each participant was equipped with a smartphone and required to install data collection software. These devices were set to be permanently discoverable by Bluetooth and scanned for nearby Bluetooth devices at 5 min intervals (hence, the detected contacts have durations which are multiples of 5 min). Proximity was assessed by measuring the received signal strength (RSSI) of Bluetooth signals from nearby devices: a high RSSI means that the two devices are expected to be physically close, a low measure indicates that devices are further apart or that there are obstacles in between, although no precise relation between RSSI and distance can be assessed in general [44]. The OD dataset was collected in an office workplace staffed with about 200 people who agreed to wear proximity sensors for several days while in the workplace [41,42]; these sensors used low-power radio communication as a proxy for the close-range face-to-face proximity of individuals wearing the devices with a temporal resolution of 20 s [32]. Both the SD and OD datasets represent proximity interactions as a temporally ordered sequence of contact networks, where nodes are

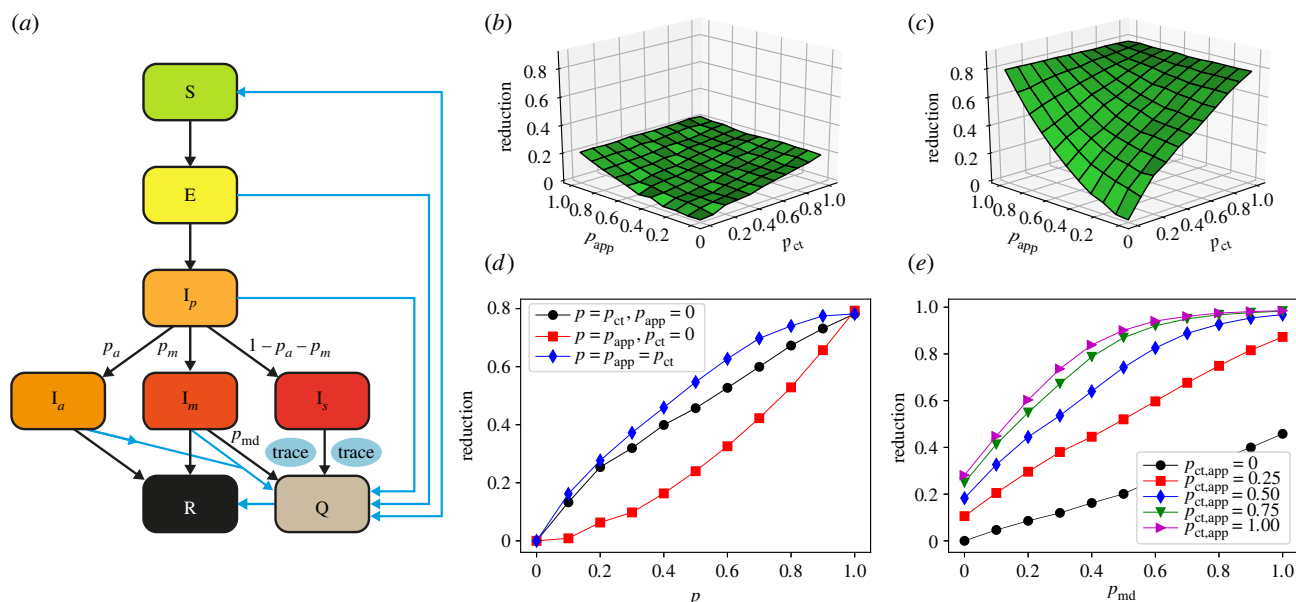


Figure 1. Model and impact of contact tracing on the epidemic size. (a) Schematic illustration of the investigated epidemic model. (b) Relative reduction in the average epidemic size as a function of MCT probability p_{ct} and app adoption rate p_{app} , with respect to the situation of only case isolation, with $p_{md} = 0.5$, for OD, with $\theta_{ct} = \theta_{app} = 15$ min. Here $\tau_{ct} = 2$ days, $\tau_{dct} = 0$. (c) The same for SD (threshold $\theta_{ct} = \theta_{app} = 15$ min). (d) Same as (c), for three slices: along the p_{ct} axis with $p_{app} = 0$, along the p_{app} axis with $p_{ct} = 0$ and along the diagonal $p_{ct} = p_{app}$. (e) Relative reduction as a function of the testing and diagnosis rate of mildly symptomatic p_{md} , for the SD with various values of $p_{ct} = p_{app}$.

individuals and an edge between nodes indicates a proximity relation. The two datasets share both similarities and differences (see also Methods): owing to the face-to-face condition, the closer range and the smaller population size for the OD, the number of contacts registered per individual is much smaller; moreover, the 5 min temporal resolution of the SD means that the shortest possible contact lasts 5 min (20 s for OD), so that the contact times are larger for the SD than for the OD data. Despite these differences, similar complex temporal and structural features are observed [35]. Moreover, the different definitions of contacts in OD and SD could correspond to simulating different spatial ranges of disease transmissions and at-risk contacts for droplet-transmitted pathogens such as severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2).

We run simulations starting from one initially infectious individual chosen at random, letting the epidemic evolve until no infectious individuals are present in the population. We first calibrate the rate of disease transmission from an infectious person to a susceptible person (see Methods) for each dataset to obtain a reproduction number $R_0 \approx 3$ in the absence of any containment measures [4,45]. Note that, because of the difference in the definition of a contact and in the temporal resolution of the two datasets, the transmission rate cannot be considered as the same for both datasets, but we instead fix R_0 as the parameter defining the effective epidemic spreading capacity. We also considered simulations with lower R_0 values, corresponding, for example, to re-emerging outbreaks during which a number of public health measures such as enhanced hygiene and mask-wearing are enforced to limit the spread, thus yielding an effectively lower R_0 : such measures are simulated by changing the transmission rate. Moreover, we considered as a possible result of behavioural change the removal of a fraction (25%) of contacts at random. As a baseline intervention, used as a reference to assess the relative improvement of contact tracing, we then considered a simple strategy that does not rely on any contact information: isolation (Q) of all cases with severe symptoms as well as isolation of a fraction p_{md} of the mildly symptomatic cases, until

they recover. For $p_{md} = 50\%$, this intervention yields a relative decrease in the average epidemic size of, respectively, 14% for the OD data and 26% for the SD data (lowering the average final epidemic size from 73% to 63% of the population in the OD and from 55% to 41% in the SD).

We then simulate MCT for all detected infectious cases (i.e. all severe cases I_s and a fraction p_{md} of the mild cases I_m): for each case, an interview is assumed to be performed to identify all individuals who have been in contact with the case for an interval longer than θ_{ct} over the 48 h before detection. To take into account the limitations of MCT, we assume that only a fraction p_{ct} of the contacts are identified, and that a delay of $\tau_{ct} = 2$ days takes place between the detection of a case and the quarantining of a recalled contact. In addition to MCT, we consider the possibility of DCT: if the identified case has a smartphone with a contact tracing app installed, a warning is sent to other app users who have been in proximity (as assessed by the contact data with the case for an interval longer than θ_{app} over the 48 h prior to detection). Since we assume that all such contacts are identified by the app, the key parameter for DCT is the fraction p_{app} of the population who has adopted the app. While we will mainly consider the value of θ_{app} set in the actual apps, we explore also larger values, which can be interpreted as what happens if a fraction of the actual contacts is not detected by the app, e.g. because of Bluetooth limitations. Moreover, we take into account potential delays due to possible app limitations (e.g. a limited number of exposure queries per day by the app) or to the time needed for an individual to self-isolate by introducing a delay τ_{dct} between warning and quarantine (τ_{dct} varying between 0 and 2 days).

The reduction in the average final epidemic size obtained by MCT and DCT, with respect to the situation of only case isolation, depends on the contact tracing parameters and efficiency. For simplicity we assume $\theta_{ct} = \theta_{app}$ as these values are fixed by public health guidelines, and we vary the fraction p_{ct} of contacts identified by MCT and the fraction of app adopters p_{app} in the population. Note that, even if p_{app} in the general population is limited by the smartphone

penetration, we consider here specific contexts where this penetration could be as high as 100% and we therefore explore the whole range of p_{app} values between 0 and 1.

Figure 1*b–d* shows how the reduction in epidemic size depends on these parameters. In particular, the effect is proportional to p_{ct} in the absence of DCT, while it grows quadratically with the fraction of app adopters in the absence of MCT. It is thus rather small at low app adoption and improving MCT leads to stronger effects. However, even in the hypothetical scenario of perfect MCT, DCT still improves the results owing to the shorter delay between detection and quarantine made possible by DCT, and the combination of DCT and MCT makes it possible to reach strong epidemic suppression effects even at intermediate parameter values.

A larger epidemic size reduction is observed for the SD case: this can be ascribed to the fact that the fraction of contacts lasting more than 15 min over 2 days (i.e. which can be traced by contact tracing) is much larger for SD (on average 59% of the individuals with whom a person has been in contact, corresponding to 96% of all the contact times) than for OD (on average 6% of the individuals with whom one has been in contact, corresponding to 45% of all the contact times). However, the qualitative behaviour of the epidemic size reduction is similar for both datasets, despite the differences in contexts and data collection methods. In the electronic supplementary material, we show that this behaviour is robust for other datasets and other parameter values (electronic supplementary material, figures S2–S8). To mimic a behavioural change, we consider in particular several reduced values of R_0 obtained by a reduction in the transmission rate. Interestingly, we observe that the relative effect is larger for smaller values of R_0 (electronic supplementary material, figure S5). Moreover, as the final epidemic size even without interventions is smaller for lower R_0 (and the epidemic is also slower), the number of cases to which contact tracing needs to be applied is smaller, and thus the cost in terms of quarantines is also smaller. We also show in the electronic supplementary material the results obtained with datasets in which a fraction of contacts has been removed (electronic supplementary material, figure S6). We moreover investigate the sensitivity with respect to an increased value of the fraction of asymptomatics (electronic supplementary material, figure S7), and to reduced compliance of the users (electronic supplementary material, figures S2–S4).

Figure 1*e* illustrates the crucial role of detecting cases with mild symptoms. Even in the absence of tracing ($p_{\text{app}} = p_{\text{ct}} = 0$), increasing the probability of case detection increases linearly the reduction in the average epidemic size. This effect becomes even stronger when contact tracing is implemented and the combination of a large probability of detection and of contact tracing can make a strong suppression of the epidemic reachable. Finally, as shown in the electronic supplementary material, delays in the warning and quarantining of individuals using the app, as modelled by τ_{dctv} , slightly decreases the efficiency of the contact tracing (electronic supplementary material, figure S8), and different values of the thresholds θ_{ct} and θ_{app} lead to qualitatively robust results (electronic supplementary material, figures S2–S4).

Contact tracing brings clear benefits as explored above, but also has a cost that can be quantified by the number of quarantine events in the population (note that a given individual might potentially be quarantined more than once). Moreover, a fraction of these quarantines are unnecessary, as

some of the quarantined individuals are, in fact, healthy. Figure 2*a–c* explores this issue. The number of quarantine events, i.e. the cost of the intervention, initially grows with p_{app} and p_{ct} as expected. However, in some cases, this cost goes through a maximum and decreases for large enough DCT app adoption or MCT efficiency. This corresponds to the fact that a very efficient contact tracing procedure can actually stop an outbreak early on, so that the epidemic does not spread and few quarantines are necessary (this is indeed the case for large p_{app} and p_{ct} in the SD case, as seen in figure 1*c*). In fact, figure 2*c* shows cost–benefit curves of the contact tracing interventions, at fixed tracing parameters, i.e. the normalized number of quarantine events as a function of the reduction in average epidemic size. These curves all exhibit a typical parabolic shape, with a cost first increasing as a function of the benefit, reaching a maximum and then decreasing. Note that each curve puts together results obtained either by MCT, by DCT or by a combination of both: in other words, the cost depends on the benefit and not on the way (MCT or DCT) this epidemic size reduction was obtained. We also note on the other hand that the cost depends on the specific value of the threshold $\theta_{\text{ct}} = \theta_{\text{app}}$ used to define at-risk contacts in the contact tracing procedure: lower thresholds make it possible to obtain higher benefit at large adoption or MCT efficiency (large p_{ct}), but the maximal value of the cost, reached at intermediate adoption or MCT efficiency, is then larger. Finally, increased delays in contact tracing lead to slightly increased costs for a given epidemic size reduction (see electronic supplementary material, figure S9). This can be expected as delays mean that some of the individuals who have been in contact have time to become infectious and transmit the disease further, leading to a larger number of quarantines. On the other hand, very large values of case detection probability p_{md} yield both higher benefits (figure 1*e*) and lower costs (see electronic supplementary material, figure S10).

Figure 2*d* shows also the behaviour of the fraction of non-infected quarantined individuals. While this fraction depends only slightly on app adoption and MCT efficiency, it increases slightly overall as the epidemic reduction increases. This fraction also increases if the delay between case detection and contact quarantining increases, and decreases if the detection probability of mild cases p_{md} increases.

Our results are obtained with simulations of a compartmental model performed on state-of-the-art datasets describing human interactions, which include many complex features of such interactions, such as heterogeneities of contact durations, intermittent formation of groups, etc. Interestingly, we can actually obtain a very similar phenomenology using analytical arguments on simplified models. To this end, we consider a Susceptible–Infectious–Recovered (SIR) model on a static random network. Using a well-known mapping of SIR models to percolation problems [46], it is indeed possible to obtain an equation for the final epidemic size of such a process. In the electronic supplementary material, we show that it is possible to extend the percolation arguments to introduce both DCT and MCT, and we show in figure 3*a* that the resulting shape of the epidemic size reduction as a function of the tracing parameters p_{ct} and p_{app} closely reflects the phenomenology observed in our numerical simulations.

Moreover, the parabolic shape of the cost–benefit curves shown in figure 2*c* can also be explained by a simple argument involving two competing effects. (i) The more we want to reduce the epidemic size, the more we need to

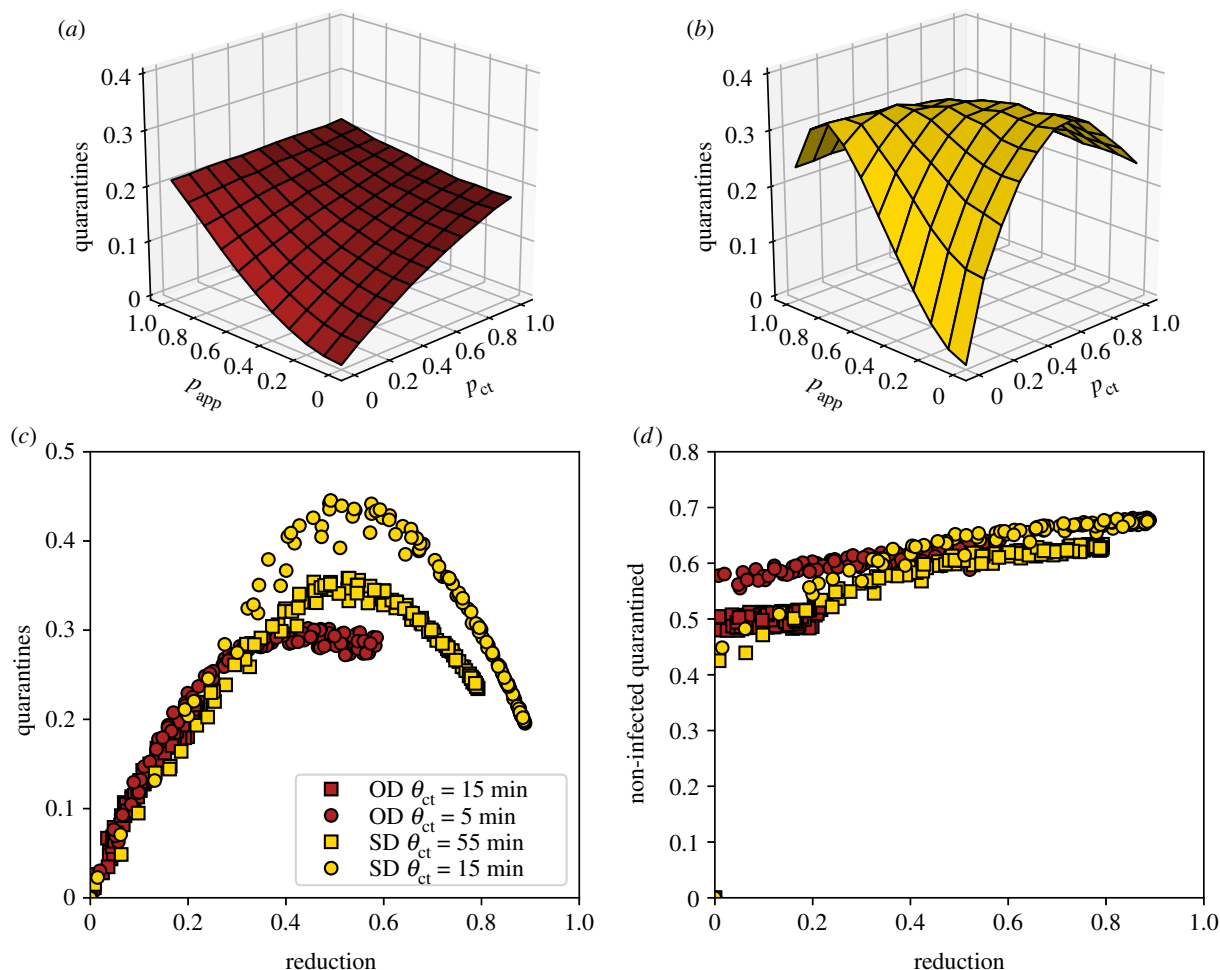


Figure 2. Number of quarantine events normalized by the population size. We show this number as a function of p_{app} and p_{ct} for (a) the OD data and (b) the SD data. (c) Number of quarantine events normalized by the population size as a function of epidemic size reduction, for the OD data (red squares; red circles for lower tracing threshold $\theta_{app} = \theta_{ct} = 5$ min), and the SD data (yellow circles; yellow squares for higher threshold $\theta_{app} = \theta_{ct} = 55$ min). (d) Ratio of non-infected quarantined individuals to quarantines, for the same datasets as in (c).

quarantine, and (ii) the more the epidemic size is reduced, the fewer quarantines are actually needed. In the absence of tracing, the epidemic size is not reduced but also no quarantines are being implemented (zero benefit and zero cost); at the other extreme the quarantines are done so aggressively that the disease can only infect a very small fraction of the population and the total number of quarantines (contacts of infected individuals) remains low as well. At the peak the quarantine rate is still low so that the disease spreads to a significant part of the population, but it is large enough so that many of the connections of the infected individuals are placed under quarantine.

The above argument can be formalized as follows. We denote the epidemic size reduction owing to contact tracing as $r = 1 - I_r/I_0$, where I_0 is the epidemic size without quarantine measures and I_r is the epidemic size with the measures in place. First, as a mean-field assumption, we express the number of quarantines q as simply the product of the average number of contacts of a person during the infectious period k , the total number of infected people I_r and the quarantine probability q_r , i.e. the probability that any contact with an infectious person leads to a quarantine: $q = q_r k I_r$. Quarantines, however, yield a decrease in the epidemic size so that the number of infectious people is itself a decreasing function of the quarantining probability: $I_r = f(q_r) I_0$, with f a decreasing function. In other terms $r = 1 - I_r/I_0$ is an increasing function of q_r . Figure 1b,c shows that the reduction in the epidemic

size, r , is roughly linear with the fraction of people asked to quarantine owing to MCT, p_{ct} . As the probability of a contact being quarantined q_r increases as well linearly on average when we ask more people to quarantine, this linearity is carried over to yield $r \propto q_r$. Inserting this as well as $I_r = (1 - r)I_0$ into the above expression for q leads to $q \propto r - r^2$, i.e. a downwards opening parabola very similar to figure 2c, as shown in figure 3b. More details and examples are provided in the electronic supplementary material.

We finally show—in figure 4—the impact of a recursive DCT. At fixed θ_{app} and app adoption, the obtained reduction in the epidemic size is larger than with the normal app. The difference relative to the usual DCT is however quadratic in p_{app} , meaning that this effect remains small unless app adoption is very high. Moreover, the resulting number of quarantine events is larger and depends more on the other parameters of the contact tracing. In fact, figure 4b,c shows that, at a given value of the epidemic size reduction, the number of quarantines is typically larger for the recursive DCT than for the usual DCT, and that the fraction of unnecessary quarantines becomes larger as well. In other words, the cost to obtain a given benefit is higher.

3. Discussion

In the fight against the COVID-19 pandemic, non-pharmaceutical interventions play a crucial role [19]. Moreover,




Figure 3. Results reproduced with simplified theoretical models. (a) Reduction in final epidemic sizes as a function of MCT probability p_{ct} and app adoption rate p_{app} . For easier comparison with the simulation results, we have matched the parameters such that the epidemic sizes at $p_{app} = p_{ct} = 0$ and $p_{app} = p_{ct} = 1$ are the same as for the SD data shown in figure 1c. The overall shape of the surface is very similar, with MCT leading to an almost linear reduction while the app adoption leads to a convex shape. (b) Number of quarantine events normalized by the population size as a function of epidemic size reduction, obtained by the simple mathematical argument explained in the main text. The height of the curve depends on the parameters and the end of the curve is determined by the maximum reduction in epidemic size. For easier comparison with figure 2c, we matched the heights and the maximum reductions to those of the simulation results.

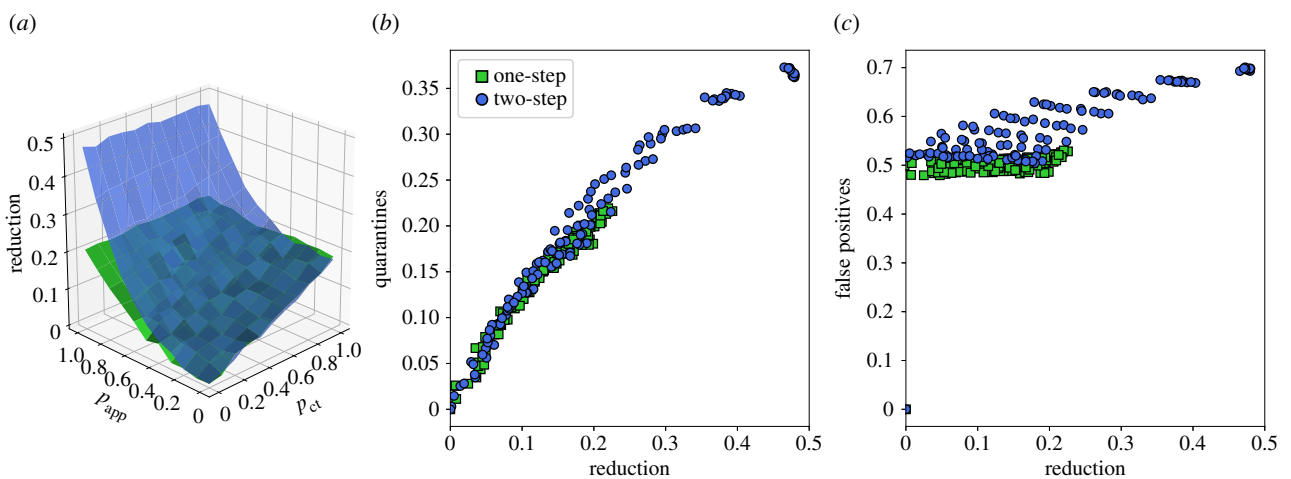


Figure 4. Recursive contact tracing versus single-step contact tracing. (a) Blue surface: epidemic size reduction as a function of MCT probability p_{ct} and app adoption rate p_{app} for recursive tracing (blue) and single-step tracing (green) for reference, for the OD dataset. $\theta_{ct} = \theta_{app} = 15$ min. (b) Normalized number of quarantine events versus epidemic size reduction for recursive ('two-step', blue) and single-step tracing (green). (c) Proportion of non-infected quarantined individuals for the same cases.

in order to limit societal costs, targeted measures such as isolation of cases, contact tracing and quarantining of contacts are considered as essential in containing potential re-emergence of outbreaks [8,13,14,16]. As traditional MCT is labour intensive and limited by people's ability to correctly remember contacts, app-based DCT is seen as a potentially useful complement and has been deployed in several countries [18,21,47]. Its actual efficiency, however, has been debated, in particular with respect to the level of adoption needed for it to make a difference [18,29,30,47,48].

In this study, we have provided new theoretical and qualitative understanding of the relative and combined efficiency of MCT and DCT by leveraging state-of-the-art datasets describing contacts between individuals in different settings, namely among office workers and among students of a university. Indeed, most previous works have considered either a homogeneous mixing hypothesis or model networks of

interactions, while real contacts are known to display a wealth of structural and temporal heterogeneities having an impact on epidemic propagation. Even if the available contact network datasets correspond to limited contexts and populations, they do encompass all these heterogeneities. Their use in simulations is therefore a crucial step to consolidate our theoretical and practical understanding of the mechanisms of various interventions and, here in particular, to understand how their impact depends on the contact tracing parameters. Here, we have considered a compartmental model drawn from the recent literature on the propagation of the SARS-CoV-2 pathogen and simulated several mitigation measures, focusing on the reduction in the final average epidemic size in the population. The isolation of severe cases and of a fraction of mildly symptomatic cases, although beneficial, is not enough to contain the spread as pre-symptomatic individuals are known to be infectious, and as a fraction of infectious are

moreover asymptomatic. We have then simulated contact tracing, both manual (MCT) and digital (DCT), and shown that both can yield a strong decrease in the final epidemic size. The relative impact of contact tracing is actually larger for lower reproduction numbers, and its cost lower: lower R_0 means indeed a slower epidemic reaching fewer individuals even without interventions, so that contact tracing needs to be applied to fewer cases, and fewer people need to be quarantined to mitigate the epidemic (this is in agreement with the results of [30]). Contact tracing thus becomes more efficient in a situation where the epidemic is already partially mitigated by other measures that keep R_0 as low as possible, showing the importance of combining contact tracing with other measures such as masks or hand hygiene, and that contact tracing might lose efficiency at too large R_0 . The epidemic size reduction moreover depends on the efficiency of the manual tracing, quantified by the fraction of contacts recorded for each detected infectious individual, and on the app adoption in the case of DCT. We find that this reduction grows linearly with the efficiency of MCT but only quadratically with the app adoption, as both case and contact need to have the app installed for a contact to be detected. We have shown that, at a qualitative level, this overall behaviour can be recovered analytically in a simplified epidemic model. We note that, as our study focuses on specific contexts (offices and university), the app adoption could potentially reach 100% in the corresponding populations, leading in that case to a very strong suppression of the epidemic. We have also shown that the cost of the intervention, as quantified by the number of quarantines, initially grows as the MCT efficiency and app adoption increase, but can become smaller if these parameters are high enough so that the epidemic is very efficiently contained. The cost–benefit curves show a typical parabola shape that can be understood using simple arguments: for low efficiencies of contact tracing, the cost increases with the reduction of epidemic size, but if the efficiency becomes large enough, the strong suppression of the outbreak leads to fewer cases, thus fewer contacts and fewer quarantines. We have directly shown how a simple analytical argument recovers this parabolic shape of the cost–benefit curves.

We also note that the DCT we have simulated does not imply knowledge of the contact network of infectious individuals, but simply that their contacts receive a warning and quarantine accordingly. This confirms that it is possible to develop strongly privacy-preserving protocols [21] that might reach large app adoption levels and thus yield a strong impact. Furthermore, we have shown that recursive DCT yields an increased impact, which, however, also grows only quadratically with app adoption. The added benefit thus remains small except at very high adoption, and comes at an increased cost in terms both of quarantine events and of non-infected quarantined individuals. Moreover, it is important to remark that recursing the contact tracing process over contacts of an index case entails building an explicit representation of the two-step networks of confirmed cases, exposing significantly more network information about those individuals than regular contact tracing. This additional network information dramatically increases privacy risks, as it can be more readily used to match the contact graph around a given user to social network data from other sources (e.g. online social networks, mobile call networks, organizational networks, etc.), increasing the probability of re-identification. The resulting privacy concerns might lead to lower app adoption in the

public, and thus potentially to an actual loss of efficiency of the DCT efforts.

Our study has a number of limitations which are important to make explicit. In particular, the data we use correspond to limited social environments (a university campus and a workplace) and we do not provide an overall general study that includes all possible multiple and differentiated contexts and their mutual interplay. These data moreover do not include interactions between the individuals participating in those studies and the general population. Nor do they include information on the precise location of the contacts, for instance whether these contacts occurred indoors or outdoors, which is known to be important in the COVID-19 context. In addition, the data were collected in pre-pandemic situations, and awareness of the spread as well as imposed restrictions have changed the behaviour of individuals in ways difficult to apprehend. Nevertheless, these datasets correspond to the current state-of-the-art in terms of data describing human interactions, and we have also considered several ways of taking into account in an effective way behavioural changes, from reduced reproduction numbers to filtering of a fraction of contacts. Moreover, one could expect that post-lockdown scenarios would rapidly lead to a return to pre-pandemic behaviour in terms of contact patterns. Overall, the datasets, even with these limitations, include complex features of real contact patterns, and it was thus important to understand the behaviour of contact tracing interventions in simulations on such data, and not only on simplified models of interactions. The results show a robust qualitative behaviour for data obtained in different contexts with different data collection infrastructures (inducing different spatial and temporal resolutions), in terms of both the epidemic size reduction and the dependency of the cost of the intervention on the MCT efficiency and on the app adoption. Moreover, our analytical approach gives us a theoretical understanding of the mechanisms behind this behaviour and shows its generality. A detailed quantification and prediction of the size of the effects would require more detailed models of specific (possibly large-scale) populations taking into account age stratification and other socio-economic factors, but important insights have already been gained by our theoretical approach.

For instance, an important point to emphasize is that, although the effect of DCT is only quadratic in the app adoption, there is no threshold effect: any increase in adoption leads to an improvement in the epidemic spread mitigation. This is actually also true for MCT, showing that the quality of interviews and any improvement in the ability to correctly find contacts of infectious cases brings an important benefit. Overall, DCT, which is also able to trace contacts between individuals who do not know each other, yields an interesting complement to MCT, and the combination of MCT and DCT is able to suppress outbreaks at limited cost if the app adoption and MCT efficiency are sufficiently high.

As the app adoption in the general population is necessarily limited by smartphone penetration, it is finally interesting to emphasize that this limitation does not necessarily hold in a specific context such as a workplace or a university. In fact, one could expect assortativity or group effects to induce a very high adoption in such specific contexts or populations. This could make it possible to very rapidly suppress outbreaks in these specific populations, at minor costs in terms of quarantines. A full investigation of such assortativity effects is an interesting direction for future work, and could

lead to policy indications on how to best focus testing and MCT resources towards populations where app adoption is expected to be lower.

4. Material and methods

4.1. Data

We use state-of-the-art publicly available datasets describing contacts between individuals in different settings, with high spatial and temporal resolution.

- The SD dataset [23,33] describes the interactions of 706 students, as registered by the exchange of Bluetooth signals between smartphones, for a period of one month. Each participant in the study was equipped with a Google Nexus 4 smartphone (used as the primary phone) and required to install study-designed data collection software. For Bluetooth measurement, all devices in the experiment were configured to be permanently discoverable, and to scan for nearby Bluetooth devices at 5 min intervals. The devices measured and recorded the RSSI: a high RSSI means that the two devices are physically close; a low measure indicates that devices are further apart or that there are obstacles in between [44].
- The OD dataset was collected by the SocioPatterns collaboration, using an infrastructure based on wearable sensors that exchange radio packets, detecting close proximity (less than or equal to 1.5 m) of individuals wearing the devices [32]. The temporal resolution of the data was 20 s. The data were collected among 232 individuals in offices during two weeks in 2015 [42]. If the epidemic simulated on the data lasts more than the dataset duration, we simply replicate it [49] by restarting from the beginning of the dataset.
- We show in addition in the electronic supplementary material results obtained using data collected among more than 300 students in a French High School during one week by the SocioPatterns collaboration, i.e. with the same infrastructure as the OD data [17].

In the SD case, the average number of distinct individuals contacted by a random individual during 1 day (average degree in the daily aggregated network) varies between 4 and 47, with strong variations between weekdays and weekends.

In the OD case, the average daily degree does not reach such large values, being between 8 and 10 depending on the day.

Moreover, the fraction of links with weight above 15 min over a period of 48 h is 59% for SD, corresponding to 96% of the total contact duration, while it is only 6%, corresponding to 45% of the total contact duration, in the OD case.

We refer to [33,42,50] for a detailed description of the datasets.

4.2. Epidemic model

We consider the compartmental epidemic model described in [4], which was designed to describe the various stages of the COVID-19 disease.

Susceptible individuals (S) can become infected upon contact with an infectious one, and then enter a latent phase (E) of duration τ_e . They then become pre-symptomatic (I_p) during τ_p . Pre-symptomatic individuals can then become either asymptomatic (I_a) or develop mild or severe symptoms, entering, respectively, the compartments I_m or I_s . This happens with probabilities a , m , $s = 1 - a - m$. Recovery occurs with rate μ .

The probability per unit time of a susceptible becoming infectious upon contact with an infectious with severe symptoms is β . On contact with a pre-symptomatic who later on develops severe symptoms, the probability per unit time becomes $r_p\beta$, and it is $r_\beta\beta$ upon contact with an asymptomatic or an infectious with mild

Table 1. Parameters of the compartmental model, taken from [4].

| parameter | value |
|-----------|----------|
| τ_e | 3.7 days |
| τ_p | 1.5 days |
| $1/\mu$ | 2.3 days |
| R_0 | 3 |
| a | 0.2 |
| m | 0.72 |
| r_β | 0.5 |
| r_p | 1 |

symptoms, or with a pre-symptomatic who then develops mild or no symptoms.

The parameter values are given in table 1. The value of β is adjusted in order to fix the basic reproduction number to a desired value in the absence of interventions. More precisely, we define R_0 as the average number of secondary infections from an initially infectious seed taken at random. We perform 4×10^3 simulations for each value of β , adjusting it until the desired value of R_0 is obtained. For $R_0 = 3$, this yields $\beta = 1.37 \times 10^{-3} \text{ s}^{-1}$ for OD and $\beta = 2.1 \times 10^{-5} \text{ s}^{-1}$ for SD. We have also considered other values of β in order to obtain lower values of R_0 corresponding to the effect of restrictions such as mask wearing or enhanced hand hygiene.

We have also considered the case $a = 0.4$, $m = 0.52$.

4.3. Interventions

We consider here a series of targeted interventions, i.e. that aim at preventing further transmissions from infected individuals.

4.3.1. Isolation

The first intervention consists simply in isolating infectious individuals once they are identified. This happens with all infected developing severe symptoms, as they will reach out to health services. Moreover, it can also happen for a fraction of infected developing mild symptoms. The probability that an individual with mild symptoms is identified as infectious is p_{md} . Moreover, we take into account that the reaction to symptoms is not instantaneous by introducing a delay $\tau_{\text{to isol}}$ between the appearance of symptoms and the isolation. Upon isolation, an infectious stops having contacts and becomes unable to transmit the disease.

Pre-symptomatics and asymptomatics cannot be identified and therefore are not isolated.

4.3.2. Manual contact tracing

When an infected individual (a 'case') is identified, they are interviewed by healthcare workers and asked to remember their contacts of the last m_{keep} days. The persons who have been in contact with the case for a cumulative time longer than θ_{ct} during these days are then warned and asked to quarantine for $\tau_q = 14$ days.

The imperfection of MCT is taken into account by the following parameters:

- only a fraction p_{ct} of the contacts longer than θ_{ct} are recalled;
- only a fraction p_c of the contacts agree to quarantine; the others do not act;
- there is a delay τ_{ct} between the detection of the case and the quarantine of their identified contacts.

The health of the quarantined individual is monitored. Therefore, if a quarantined person develops symptoms (even mild),

Table 2. Parameters of the interventions. We also explore $p_c = 0.6$ and the whole range $0 \leq p_{md} \leq 1$.

| parameter | value |
|-------------------|------------------|
| $\tau_{to\ isol}$ | 0.5 day |
| p_{md} | 0.5 |
| τ_q | 14 days |
| n_{keep} | 2 days |
| p_c | 1 |
| τ_{ct} | 2 days |
| τ_{dct} | from 0 to 2 days |

they are interviewed to trace their recent contacts, with the same procedure as the initial case. The parameter values are given in table 2.

4.3.3. Digital contact tracing

We consider, either by itself or in addition to MCT, the possibility that individuals have installed a privacy-preserving proximity-tracing app. The app registers proximity events with other individuals equipped with the app.

Whenever an app-adopter is diagnosed as infected (the case), the anonymous random identifiers used by their app in the last n_{keep} days are uploaded to the central server (see [21] for details).

All app-adopters regularly check the central server and their app compares the list of identifiers of infected individuals with the list of identifiers received in the previous days. If the app detects that the cumulated time in contact with (one or several) infected app-adopters in the last n_{keep} days exceeds θ_{app} , it triggers a warning to contact health authorities and go into quarantine.

We assume that the transmission of information is instantaneous, but a delay τ_{dct} can take place between the warning of an individual by the app and the start of their quarantine. The other parameters are the fraction of app-adopters in the population, p_{app} , and the probability p_c that an individual receiving a warning goes into quarantine. Non-compliant individuals do not change their behaviour. If a quarantined app-adopter develops symptoms (even mild), the process is iterated, i.e. their identifiers are uploaded. The parameter values are given in table 2.

References

- COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) at Johns Hopkins University. See <https://github.com/CSSEGISandData/COVID-19>.
- Guzzetta G *et al.* 2020 The impact of a nation-wide lockdown on COVID-19 transmissibility in Italy. (<http://arxiv.org/abs/2004.12338>).
- Zhang J *et al.* 2020 Changes in contact patterns shape the dynamics of the COVID-19 outbreak in China. *Science* **368**, 1481–1486. (doi:10.1126/science.abb8001)
- Domenico LDJ, Pullano G, Sabbatini CE, Boëlle P-Y, Colizza V. 2020 Impact of lockdown on COVID-19 epidemic in Île-de-France and possible exit strategies. *BMC Med.* **18**, 240. (doi:10.1186/s12916-020-01698-4)
- Koo JR, Cook AR, Park M, Sun Y, Sun H, Lim JT, Tam C, Dickens BL. 2020 Interventions to mitigate early spread of SARS-CoV-2 in Singapore: a modelling study. *Lancet Infect. Dis.* **20**, 678–688. (doi:10.1016/S1473-3099(20)30162-6)
- Hsiang S *et al.* 2020 The effect of large-scale anti-contagion policies on the COVID-19 pandemic. *Nature* **584**, 262–267. (doi:10.1038/s41586-020-2404-8)
- Carrat F *et al.* 2020 Seroprevalence of SARS-CoV-2 among adults in three regions of France following the lockdown and associated risk factors: a multicohort study. *medRxiv.* (doi:10.1101/2020.09.16.20195693)
- Ferguson N *et al.* 2020 Report 9: impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Technical report, Imperial College London, London, UK.
- Ma S, Zhang J, Zeng M, Yun Q, Guo W, Zheng Y, Zhao S, Wang MH, Yang Z. 2020 Epidemiological parameters of coronavirus disease 2019: a pooled analysis of publicly reported individual data of 1155 cases from seven countries. *medRxiv.* (doi:10.1101/2020.03.21.20040329)

4.3.4. Recursive contact tracing

We finally implement recursive (two-step) warnings for app-adopters: in this case, we assume that a protocol is in place so that app-adopters can compute the cumulative contact time during the previous n_{keep} both with app-adopting infected individuals and with app-adopting contacts of app-adopting infected individuals who have received a warning by the app. If this cumulative time exceeds θ_{app} , it triggers a warning to contact health authorities and go into quarantine.

The parameters are the same as for the DCT: the fraction of app-adopters p_{app} , and the probability p_c that an individual receiving a warning goes into quarantine. Non-compliant individuals do not change their behaviour.

4.4. Quantification of the impact of interventions

For each scenario, defined by a given set of interventions with a given set of parameter values, we perform 10 000 numerical simulations with a single, randomly chosen individual in the latent phase at the initial time. Simulations are run until no infected individuals are present in the population, i.e. all individuals are either susceptible or recovered.

In order to quantify the impact of the spread on the population, and the effect of the various interventions, we measure:

- the average final epidemic size, i.e. the fraction of individuals who are recovered at the end of the simulation;
- the fraction of population going through quarantine;
- the fraction of the quarantined individuals who were in fact susceptible.

In particular, we quantify the effect of interventions by the relative reduction in the epidemic size, defined as $\frac{I_0 - I_r}{I_0}$, where I_0 is the average final epidemic size without interventions and I_r is the average final size in the presence of interventions.

Data accessibility. We use only publicly available data. In addition, we have made our code available at https://github.com/jsaramak/contact_tracing and <https://github.com/barrat/digitalcontacttracing>.

Authors' contributions. All authors conceived the study. A.B., J.S. and M.K. performed the numerical simulations and analysis. All authors wrote the manuscript.

Competing interests. We declare we have no competing interests.

Funding. C.C. acknowledges partial support from the Lagrange Project of the ISI Foundation funded by the CRT Foundation, and from the European Union's Horizon 2020 research and innovation programme under grant agreement nos. 101003688 (EpiPose) and 101016233 (PERISCOPE). A.B. acknowledges partial support from the ANR project DATAREDEX (ANR-19-CE46-0008-01).

10. Ganyani T, Kremer C, Chen D, Torneri A, Faes C, Wallinga J, Hens N. 2020 Estimating the generation interval for COVID-19 based on symptom onset data. *Eurosurveillance* **25**, 23–30. (doi:10.2807/1560-7917.ES.2020.25.17.2000257)
11. He X *et al.* 2020 Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nat. Med.* **26**, 672–675.
12. Lavezzo E *et al.* 2020 Suppression of a SARS-CoV-2 outbreak in the Italian municipality of Vo'. *Nature* **584**, 425–429. (doi:10.1038/s41586-020-2488-1)
13. Cheng H-Y, Jian S-W, Liu D-P, Ng T-C, Huang W-T, Lin H-H. 2020 Contact tracing assessment of COVID-19 transmission dynamics in Taiwan and risk at different exposure periods before and after symptom onset. *JAMA Intern. Med.* **180**, 1156–1163. (doi:10.1001/jamainternmed.2020.2020)
14. Pinotti F, Di Domenico L, Ortega E, Mancastropa M, Pullano G, Valdano E, Boëlle P-Y, Poletto C, Colizza V. 2020 Tracing and analysis of 288 early SARS-Cov-2 infections outside China: a modeling study. *PLoS Med.* **17**, 1–13. (doi:10.1371/journal.pmed.1003193)
15. Klinckenberg D, Fraser C, Heesterbeek H. 2006 The effectiveness of contact tracing in emerging epidemics. *PLoS ONE* **1**, 1–7. (doi:10.1371/journal.pone.0000012)
16. Hellewell J *et al.* 2020 Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *Lancet Global Health* **8**, e488–e496. (doi:10.1016/S2214-109X(20)30074-7)
17. Mastrandrea R, Fournet J, Barrat A. 2015 Contact patterns in a high school: a comparison between data collected using wearable sensors, contact diaries and friendship surveys. *PLoS ONE* **10**, 1–26. (doi:10.1371/journal.pone.0136497)
18. Ferretti L, Wymant C, Kendall M, Zhao L, Nurtay A, Abeler-Dörner L, Parker M, Bonsall D, Fraser C. 2020 Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* **368**, eabb6936. (doi:10.1126/science.abb6936)
19. Kucharski AJ *et al.* 2020 Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of SARS-CoV-2 in different settings: a mathematical modelling study. *Lancet Infect. Dis.* **20**, 1151–1160. (doi:10.1016/S1473-3099(20)30457-6)
20. Firth JA, Hellewell J, Klepac P, Kissler SM, Kucharski AJ, Spurgin LG. 2020 Combining fine-scale social contact data with epidemic modelling reveals interactions between contact tracing, quarantine, testing and physical distancing for controlling COVID-19. *medRxiv*. (doi:10.1101/2020.05.26.20113720)
21. Troncoso C *et al.* 2020 Decentralized privacy-preserving proximity tracing. (<http://arxiv.org/abs/2005.12273>).
22. Bradshaw WJ, Alley EC, Huggins JH, Lloyd AL, Esvelt KM. 2020 Bidirectional contact tracing is required for reliable COVID-19 control. *Nat. Commun.* **12**, 232. (doi:10.1038/s41467-020-20325-7)
23. Stopczynski A, Sekara V, Sapiezynski P, Cuttone A, Madsen MM, Larsen JE, Lehmann S. 2014 Measuring large-scale social networks with high resolution. *PLoS ONE* **9**, 1–24.
24. Kojaku S, Hébert-Dufresne L, Ahn Y-Y. 2020 The effectiveness of contact tracing in heterogeneous networks. (<http://arxiv.org/abs/2005.02362>).
25. Kaptchuk G, Goldstein DG, Hargittai E, Hofman J, Redmiles EM. 2020 How good is good enough for COVID19 apps? The influence of benefits, accuracy, and privacy on willingness to adopt. (<http://arxiv.org/abs/2005.04343>).
26. Lorch L, Trouleau W, Tsirtsis S, Szanto A, Schölkopf B, Gomez-Rodriguez M. 2020 A spatiotemporal epidemic model to quantify the effects of contact tracing, testing, and containment. (<http://arxiv.org/abs/2004.07641>).
27. Hinch R *et al.* 2020 Effective configurations of a digital contact tracing app: a report to NHSX. See https://cdn.theconversation.com/static_files/files/1009/Report_-_Effective_App_Configurations.pdf?1587531217.
28. Kretzschmar ME, Rozhnova G, Bootsma MJC, van Boven M, van de Wijgert JHHM, Bonten MJM. 2020 Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. *Lancet Public Health* **5**, e452–e459. (doi:10.1016/S2468-2667(20)30157-2)
29. Cencetti G, Santin G, Longa A, Pigani E, Barrat A, Cattuto C, Lehmann S, Salathé M, Lepri B. 2021 Digital proximity tracing on empirical contact networks for pandemic control. *Nat. Commun.* **12**, 1655. (doi:10.1038/s41467-021-21809-w)
30. Moreno López JA, Arregui García B, Bentkowski P, Bioglio L, Pinotti F, Boëlle P-Y, Barrat A, Colizza V, Poletto C. 2021 Anatomy of digital contact tracing: role of age, transmission setting, adoption and case detection. *Sci. Adv.* **7**, eabd8750. (doi:10.1126/sciadv.abd8750)
31. Fraser C, Riley S, Anderson RM, Ferguson NM. 2004 Factors that make an infectious disease outbreak controllable. *Proc. Natl Acad. Sci. USA* **101**, 6146–6151. (doi:10.1073/pnas.0307506101)
32. Cattuto C, Van den Broeck W, Barrat A, Colizza V, Pinton J-F, Vespignani A. 2010 Dynamics of person-to-person interactions from distributed RFID sensor networks. *PLoS ONE* **5**, 1–9. (doi:10.1371/journal.pone.0011596)
33. Sapiezynski P, Stopczynski A, Lassen DD, Lehmann S. 2019 Interaction data from the Copenhagen Networks Study. *Sci. Data* **6**, 315. (doi:10.1038/s41597-019-0325-x)
34. Gauvin L, Panisson A, Cattuto C. 2014 Detecting the community structure and activity patterns of temporal networks: a non-negative tensor factorization approach. *PLoS ONE* **9**, e86028. (doi:10.1371/journal.pone.0086028)
35. Aslak U, Rosvall M, Lehmann S. 2018 Constrained information flows in temporal networks reveal intermittent communities. *Phys. Rev. E* **97**, 062312. (doi:10.1103/PhysRevE.97.062312)
36. Sekara V, Stopczynski A, Lehmann S. 2016 Fundamental structures of dynamic social networks. *Proc. Natl Acad. Sci. USA* **113**, 9977–9982. (doi:10.1073/pnas.1602803113)
37. Barrat A, Barthelemy M, Vespignani A. 2008 *Dynamical processes on complex networks*. Cambridge, UK: Cambridge University Press.
38. Masuda N, Holme P (eds). 2017 *Temporal network epidemiology*. Singapore: Springer.
39. Stehlé J *et al.* 2011 Simulation of an SEIR infectious disease model on the dynamic contact network of conference attendees. *BMC Med.* **9**, 87. (doi:10.1186/1741-7015-9-87)
40. Géniois M, Vestergaard C, Cattuto C, Barrat A. 2015 Compensating for population sampling in simulations of epidemic spread on temporal contact networks. *Nat. Commun.* **8**, 8860. (doi:10.1038/ncomms9860)
41. SocioPatterns collaboration. See www.sociopatterns.org (accessed 1 December 2020).
42. Géniois M, Barrat A. 2018 Can co-location be used as a proxy for face-to-face contacts? *EPJ Data Sci.* **7**, 11. (doi:10.1140/epjds/s13688-018-0140-1)
43. Funk S, Salathé M, Jansen VAA. 2010 Modelling the influence of human behaviour on the spread of infectious diseases: a review. *J. R. Soc. Interface* **7**, 1247–1256. (doi:10.1098/rsif.2010.0142)
44. Sekara V, Lehmann S. 2014 The strength of friendship ties in proximity sensor data. *PLoS ONE* **9**, e100915. (doi:10.1371/journal.pone.0100915)
45. Salje H *et al.* 2020 Estimating the burden of SARS-CoV-2 in France. *Science* **369**, 208–211. (doi:10.1126/science.abc3517)
46. Newman MEJ. 2002 Spread of epidemic disease on networks. *Phys. Rev. E* **66**, 016128. (doi:10.1103/PhysRevE.66.016128)
47. Colizza V *et al.* 2021 Time to evaluate COVID-19 contact-tracing apps. *Nat. Med.* **27**, 361–362. (doi:10.1038/s41591-021-01236-6)
48. Salathé M *et al.* 2020 Early evidence of effectiveness of digital contact tracing for Sars-Cov-2 in Switzerland. *Swiss Med. Wkly.* **150**, w20457.
49. Machens A, Gesualdo F, Rizzo C, Tozzi AE, Barrat A, Cattuto C. 2013 An infectious disease model on empirical networks of human contact: bridging the gap between dynamic network data and contact matrices. *BMC Infect. Dis.* **13**, 185. (doi:10.1186/1471-2334-13-185)
50. Mones E, Stopczynski A, Pentland AS, Hupert N, Lehmann S. 2018 Optimizing targeted vaccination across cyber-physical networks: an empirically based mathematical simulation study. *J. R. Soc. Interface* **15**, 20170783. (doi:10.1098/rsif.2017.0783)