Added for publication: This paper was reviewed by Graeme Ackland and James Ackland on behalf of the UKRI-RAMP initiative

Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing.

Review by Graeme Ackland ([redacted]) and James Ackland
[redacted]

The paper assesses the potential value of contact tracing in the fight against COVID. It envisages this being done via mobile app and bluetooth. It does not make any attempt to study whether the hardware or software exists or is viable.

It has gone through the review process at Science, so I will assume it meets their requirements of correctness (i.e. I didn't code it up) and general interest (obviously).

The study is based around the parameter *r*, the measure of "infection rate" which forms a key feature of epidemic modelling. This measures "how many people are infected by an infected person?", and it is dependent on the nature of the disease, and the behaviour of people. It does **not** measure "what is the chance that an encounter leads to an infection?", which depends mainly on the disease and is harder to mitigate.

From evidence about *r*, the paper back deduced a contact network and the relative importance of various ways the virus is spread, namely:

I. *Symptomatic transmission:* direct transmission from a symptomatic individual.
II. *Pre-symptomatic transmission:* transmission that occurs before the source individual experiences noticeable symptoms.
III. *Asymptomatic transmission:* transmission from individuals who never experience noticeable symptoms.
IV. *Environmental transmission:* transmission from individuals who never meet.

A phone app. can immediately track type I (but the symptomatic user is clearly non-compliant with quarantine), types III and IV are invisible to the app. The paper rigorously derives the importance of these from another model, based on *r*. This model has 7 parameters, three of which can be directly measured with high accuracy, four deduced indirectly or by anecdote. The accuracy of the model, particularly type III and IV transmission, appears to be sensitive to the choice of these parameters. According to the model, approximately 45% of the transmissions of are type II, for which the app is useful.

One could at this stage quarantine all contacts of an infected person. Doing so would reduce *r*, by 45%: useful, but not sufficiently good to bring its "natural" level of R0 to below 1. So the conclusion is that contact tracing alone would not stop the epidemic spreading through the entire population causing the widely-reported 0.5M deaths, albeit on a longer timescale. What the authors do not discuss is the number of people who would be quarantined by the app. This appears to be huge – one infected person would drive into quarantine everyone they encounter (see below for alternative approaches).

The problem then is what to do with the data, and here again we need to disentangle two things: given that someone is diagnosed with COVID, the app is likely to present a reliable way of tracing connections backward in time, with very few false positives given low COVID incidence. Conversely, going forward in time the vast majority of people will be quarantined unnecessarily – presumably everyone who comes within 2m of an infected person. The paper does not estimate the numbers going into quarantine.

Effective Configurations of a Digital Contact Tracing App: A report to NHSX
This paper assesses the effect of the app. They are more specific about the UK context. They also assume a continuing lockdown of the vulnerable over-70 group. The approximate benefit of the app is that 14% of the population are spared lockdown – which seems reasonable.
A claim is made that the app. saves lives according to Fig 4. This seems to be incorrect, as the graphs are truncated when all scenarios show rising infection rates except for no intervention, which has reached the 45M cases required for herd immunity. I believe the other graphs will also have to reach that level unless COVID is wiped out not only in the UK, but everywhere (or surveillance is continued successfully forever). So far from saving lives, the model epidemic is prolonged (the graphs may be abruptly truncated by a vaccine, but this is not in the model). The level of quarantining implied by each scenario is shown. The results here seem compatible with the Science paper, with differences that are understandable from the different parameterizations.
A careful and detailed parameter analysis (Appendices) shows the model gives very different results for plausible assumptions. This model is **not** a weather-like chaotic system: these uncertainties could be eliminated if better data became available.

Defining an epidemiologically meaningful contact from phone proximity events: uses for digital contact tracing
Addresses the question of "who to isolate". They introduce sensible heuristics based on time and proximity of contacts: these require much more exposure than the current "2m rule" implies. Phone proximity does not account for physical barriers to the virus, such as floors and walls within buildings - people self-isolating in adjacent flats appear as contacts. No estimate of the fraction of false positives flagged for quarantine is given: this can readily be estimated as (number of contacts – R0)/(number of contacts). In normal life with a "2m rule" this would be unacceptably high – getting worse if R0 drops. Using the more stringent contact criteria envisaged here will mean some transmission events are missed (insufficient data exists to determine how many).

Proximity Tracing App Report from the Measurement Campaign 2020-04-09
I didn't understand the detail of this paper, they seem to be testing how reliably "true contacts" are correlated with "Bluetooth contacts", and the answer seems to be about 70%. This would put a limit on the effectiveness of the app, and could be combined with non-compliance in the first two papers to get an estimate of the effectiveness of the app. The report compares COV and Bluetooth transmission in air, which is reasonable. I note that Bluetooth passes readily through walls, but COV does not. For indoor contacts, reliability will be much lower.

Summary

There is nothing technically wrong with these papers, but, in my opinion, they avoid discussion of the downsides of the app. Notably, the number of people released from lockdown is rather small. The level of compliance required and assumed seems to be optimistic without enforcement – in the scenarios investigated app will typically be telling people with ~5% chance of being infected to quarantine themselves (another way of saying this - they undergo 20 unnecessary quarantines on average before being infected). No special treatment is made for "essential workers" in these models, which if allowed would again reduce the effectiveness. Countries which implemented such apps. have been more successful in suppressing the disease, so I was surprised at the conclusion I came to, however those countries also implemented higher rates of testing and strong lockdown policies, so evidence for a causal connection with phone use is weak.

The work suggests that to be successful, the app users need to
1/ Achieve take up of the app at 75% (across windows, android and ios systems)
2/ Define "contact criteria" which detect at least 85% of infection interactions
3/ Use "contact criteria" leading to 12 contacts per person (i.e much more stringent than the than the widely-publicised "2-metre" rule).
4/ 80% self-quarantining prior to symptoms when advise by app.

These are important considerations, and it should be considered whether they can be achieved in practice. While clearly important, this is not my expertise. Given the short timescales, I took the liberty of asking my son to comment – he works in political psychology research, so I have some reason to value his judgment.

The clearest behavioural barrier is popular voluntary participation. Helpfully, we are in a moment in which individuals feel a heightened sense of social responsibility and in-group loyalty (e.g. NHS clap, spike in volunteering, community support networks etc.). Key to achieving motivation to download the app will be a tangible benefit to those who have it. The most obvious such benefit is the availability of information that can help individuals protect themselves and others, both in the form of contact tracing, and more generic COVID advice (as per the King's College 'JoinZoe Symptom Tracker' App). However, as with home genealogy testing services, some will take the opinion that 'if they've got it, it's too late and they would rather not know', and others are already trying to limit their exposure to COVID-related updates. Furthermore, should the app gain a reputation for quarantining false positives, it seems that the benefit to any given individual of app ownership or use will rapidly wane.

Even if the requisite 75% of people are told the app exists and are motivated to sign up, there remain the obstacles of 'computer literacy' and the technical complications that might arise from people using outdated operating systems or lacking phone storage capacity to download new apps. Further, the technology will presumably require individuals to take their phones with them every time they leave their homes, to remember to switch on Bluetooth pairing, and indeed for their phones to remain switched on with sufficient battery life. In general, small amounts of friction to behaviour change can quickly override the initial motivation to act, and there are enough such barriers that a 75% rate of downloads, let alone subsequent usage, will be difficult to achieve. By way of simple comparison, a 75% reach among all Android users would make the app the most popular user-installed app, ahead of Facebook, and with triple the popularity of the most popular game (Subway Surfers).

Another problem is that the contact criterion used by the app will necessarily be different from the “2-metre rule”. Should this discrepancy become public knowledge, it is likely to undermine confidence in the accuracy of either the public health message, the app, or both. Should the much anticipated ‘lockdown fatigue’ begin to arise, we can expect that such mixed messaging will be used as an excuse by individuals to ignore government advice.

It is clear from the section ‘ethical implications’ that the authors have taken every effort to ensure their app will be implemented responsibly and transparently. Nonetheless, there will be at least some negative reaction from those in the public and press who misunderstand its function. While it may be possible to minimise this in the mainstream, it is nonetheless clear from the recent 5G controversy that conspiracy theories can create a genuine threat to public trust. It would therefore be wise to use communications strategies to explicitly link the app to an institution with high prior public trust – the NHS the obvious candidate - so as to minimise the threat of its being undermined. Such public trust in the app is clearly a vital precursor to the engagement demanded by the model assumptions.