Contact tracing and testing – some counterfactuals

Commissioned for Channel 4 Research

Karl Friston

Dynamic causal modelling

What follows are simulations of what might have happened – and what could happen – under different testing and contact tracing scenarios. These quantitative analyses are based upon dynamic causal models of viral spread and behavioural responses, i.e., social distancing and lockdown. This kind of modelling encompasses all the factors that influence community transmission. As such, its predictions are about mitigated outcomes. This can be contrasted with conventional epidemiological forecasts that generally consider unmitigated outcomes, in the absence of interventions or behavioural responses.

The first step is to find the best model of morbidity and mortality. In other words, find the most accurate but simple explanation for the data at hand. Figure 1 shows the outcomes (black dots) used to optimise the model and its predictions over the entire outbreak in the United Kingdom. The predictions are shown as lines with shaded confidence intervals. Crucially, these predictions are of the past and future, meaning that they can be interpreted as projections[1].

In this example, seven sorts of data were used to inform the predictions; namely, the number of new PCR tests reported every day, daily COVID-related deaths, the number of COVID patients requiring mechanical ventilation in hospital, the number of PCR tests performed, the prevalence of infection (as estimated by community surveillance studies), seropositivity (as estimated from Ab tests) and the prevalence of systems (as estimated from the COVID Symptom Tracker).

[1] The appendix provides the underlying trajectories of latent states generating these outcomes (namely, location, infection, clinical and testing states).

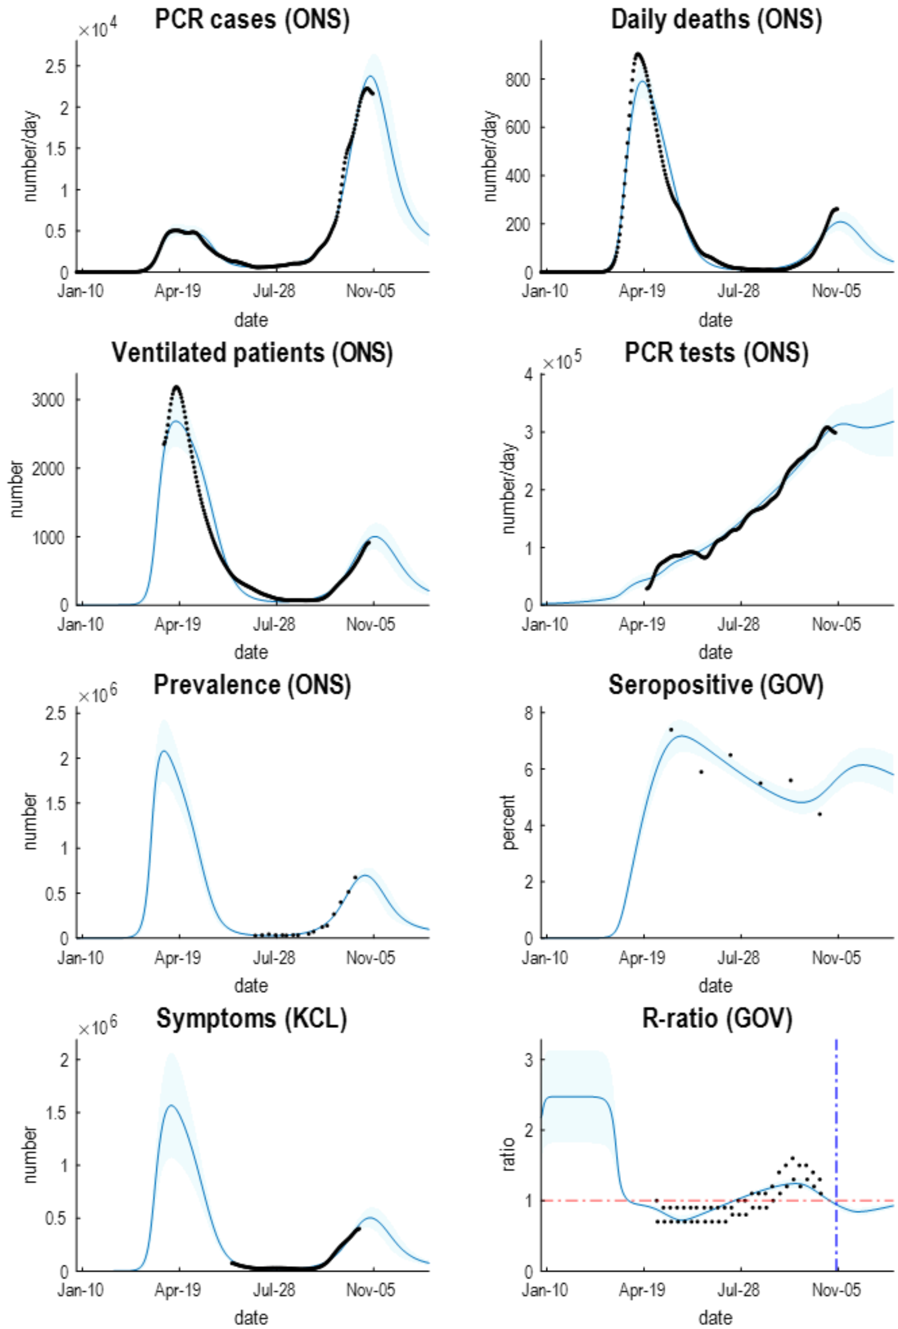

Figure 1: *predicted outcomes and 90% Bayesian credible intervals for the five kinds of data used to optimise the model parameters. All timeseries data reported in days have been smoothed to produce a seven-day average. Please see appendix for data sources and modelling details.*

It can be seen that the dynamic causal model provides a fairly accurate account of these data. Because the model includes responses to the prevalence of infection it can use implicit responses to the first wave to predict how we will respond to the secondary wave. These predictions of mitigated responses are the most likely outcomes under this coarse-grained modelling at the population level. These predictions could be regarded as idealistic; in the sense that the model does not include seasonal variations or a detailed consideration of how different communities – or age groups – intermix. However, equipped with this kind of model, we can ask what would have happened under different regimes of testing and tracing.

The relationship between testing and contact tracing is not straightforward. For example, PCR testing could have two agendas. First, it could be used to **identify cases** who are infected – so that they can be isolated. Second, one might then **trace the contacts** of people who have been identified, so that they can be supported in self-isolation. This means that there may be a direct effect of increasing testing rates and an indirect effect mediated by contact tracing. The indirect effect clearly depends upon the efficacy of contact tracing and the probability that contacts will self-isolate. In what follows, we will look at the direct and indirect effects separately and the implications of the first for the second.

Testing

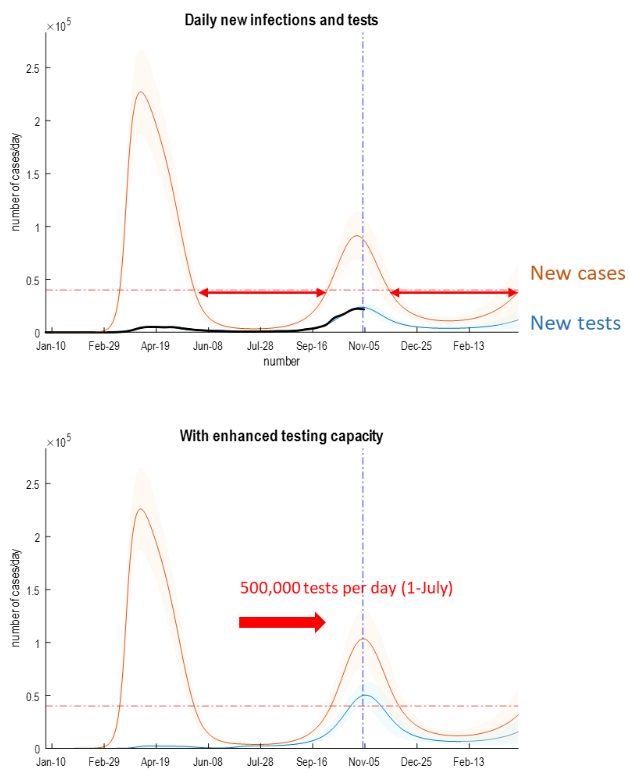

Figure 2. ***upper panel****: a comparison of the number of new PCR positive cases identified every day (cyan), in relation to the number of people becoming infected (orange), i.e., incidence. The blue lines indicate the date at which the simulations were performed, while the red line is a threshold on the incidence of infection – at 20,000 a day. The* ***lower panel*** *shows the same results but with an increase in testing capacity to 500,000 tests per day on 1 August 2020, over a period of two weeks.*

Figure 2 shows the number of PCR tests reported each day together with the number of new infections as estimated by the model. The first thing to note is that only a small proportion of new cases are picked up by testing (about 25%, depending upon the prevalence of infection and the bias towards testing people who are symptomatic). This has two implications. First, it means that simply increasing testing capacity will only catch a small proportion of new cases and will therefore have a limited direct effect on community transmission. This is illustrated in the lower panel, where we have simulated a substantial increase in testing capacity (to around 500,000 tests per day) on 1 July 2020. In the absence of any enhanced contact tracing, this has the effect of postponing the secondary wave by a few weeks but does not have a material effect on overall infection rates.

The second key observation here is that there are periods or windows of opportunity during which an effective contact tracing scheme could be pursued. These windows are when the incidence of new cases is sufficiently small to allow their contacts to be traced and supported in isolation. The estimated efficacy of contact tracing and isolation is about 4%. In other words, the probability that you will be contacted if infected but asymptomatic – and will self-isolate – is about 4%.

So, what levels of contact tracing could one reasonably hope to achieve? This depends upon the ability to trace the contacts of newly cases. If we assume that there were 10,000 contact tracers for the UK and each could handle two new cases every day, this places an upper bound of 20,000 new cases per day, over which contact tracing would be overwhelmed. The red lines in Figure 2 shows this threshold and the implicit windows of opportunity (red arrows), which disappear around the 15th of September and reappear in late November.

What would happen if we were able to increase the efficacy of contact tracing and isolation from 4% to 25%? And is this reasonable? If we assume – in the best-case scenario – that a combination of PCR testing and clinical surveillance identified about 30% of new cases and 80% of their contacts were traced and isolated, then this would result in an efficacy of 25%. Figure 3 shows the effect on fatality, if the efficacy of contact tracing was increased to 25% on 1 August (just after the point of lowest prevalence during the summer), before the secondary wave on 1 September and after the secondary wave, on 1 December.

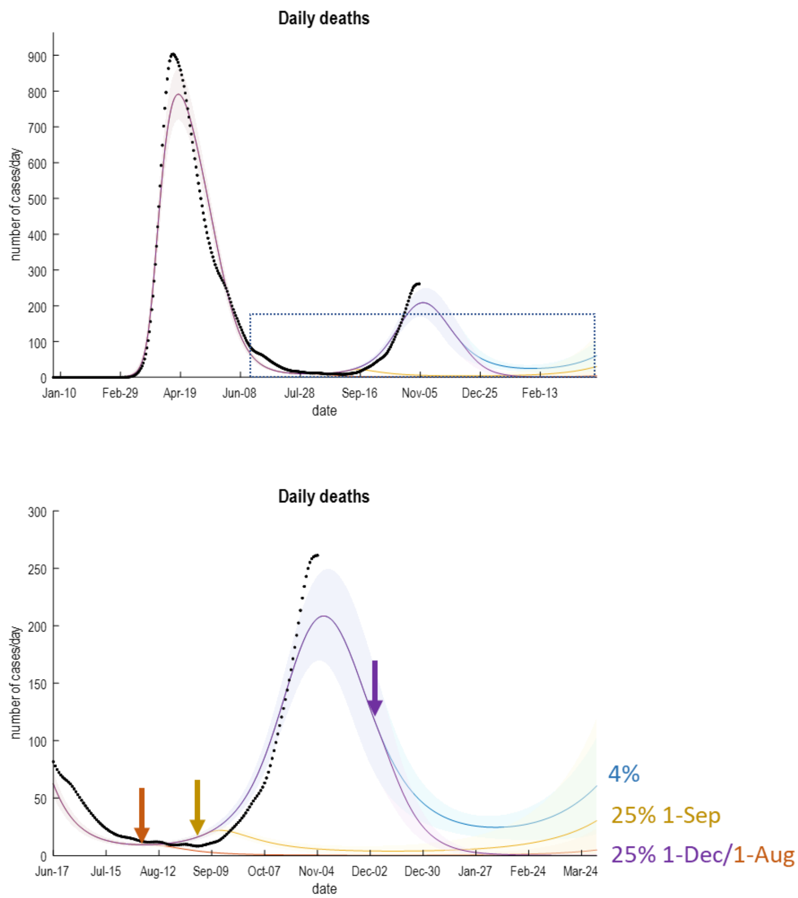

Figure 3: *the effects of enhanced contact tracing on death rates. These are indicated by the different coloured lines (and confidence intervals). They report the effects of increasing the efficacy of contact tracing from 4% to 25% on 1 August, 1 September, and 1 December 2020. The lower panel shows the same data but focusing on the current period. These data are predictions of daily deaths with reported deaths at the time of writing (24th of October) shown as black dots.*

These simulations suggest that enhanced contact tracing in late summer—at the time when the prevalence of infection was lowest—might have suppressed viral transmission and eluded a secondary way. The same enhancement a month later, in September, would have substantially mitigated the secondary wave of fatalities; however, it would only have deferred the morbidity burden until early spring.

Conversely, introducing an efficient contact tracing program in December could suppress community transmission—despite the fact the prevalence of infection will be higher than in September. One may ask why this is the case. The key difference between states of affairs before and after the secondary wave explains the differential impact of enhanced contact tracing. These differences include a small increase in population immunity that attenuates viral spread following the secondary wave (see appendix).

This analysis—with the above caveats—speak to a pressing need to exploit our increasing testing capacity over the next few weeks by increasing the efficacy of contact tracing. In principle, under the idealised assumptions above, this could suppress and possibly eliminate community transmission by next spring.

Appendix

Graphical presentation of the latent states generating the predicted outcomes in Figure 1.

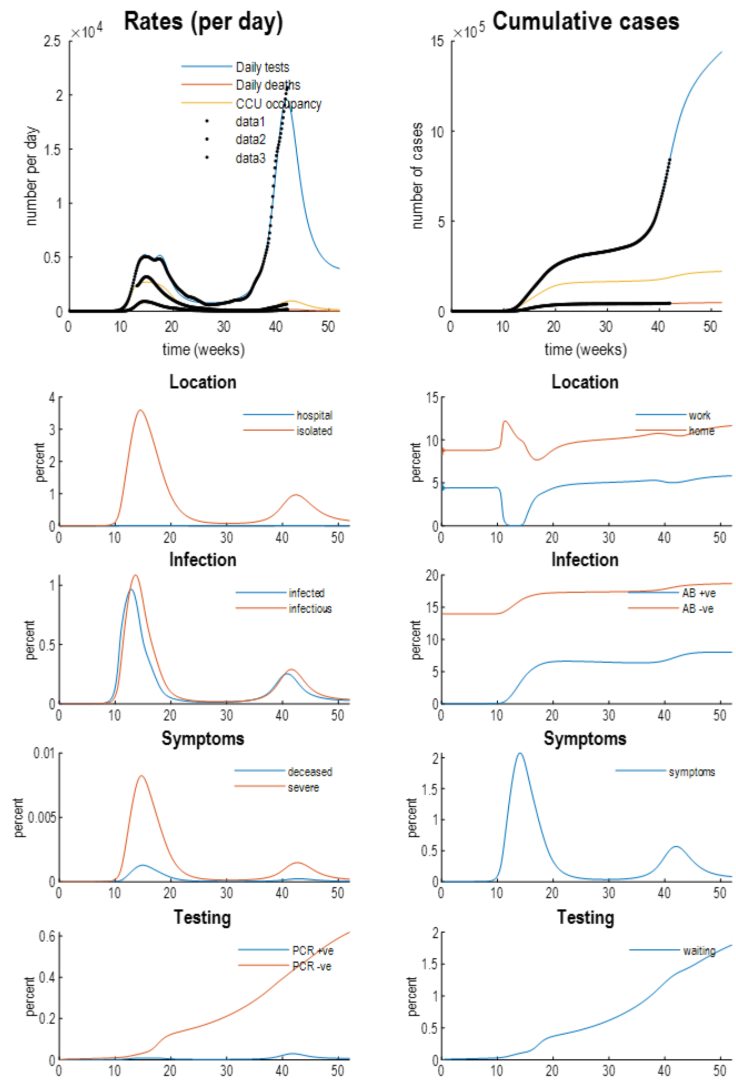

Data sources:

https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/datasets/coronaviruscovid19infectionsurveydata
https://www.ndm.ox.ac.uk/covid-19/covid-19-infection-survey/results
https://coronavirus.data.gov.uk/
https://covid.joinzoe.com/data#levels-over-time

Modelling sources and resources: https://www.fil.ion.ucl.ac.uk/spm/covid-19/