Policy variables for controlling the pandemic

Lockdowns and pandemic control - do they reduce overall casualties?

BBC Radio4Today (Mon 8th Oct) included an interview with Prof Graeme Ackland from Edinburgh University, luridly trailed as “a strict lockdown results in more deaths over the entire course of the pandemic” in a re-analysis of Prof Ferguson’s Imperial College model whose conclusions triggered the lockdown in March:

What did not come out in the interview was the rationale behind the headline. Prof Ackland said that without the lockdown the health service would have been overwhelmed with horrendous numbers of deaths.

So let’s examine that headline assertion.

What does the modelling say?

Ackland’s team checked the predictions against what happened by replicating the Imperial College “Report 9” model (with Imperial College’s assistance) on an open source model called CovidSim. CovidSim contains large amounts of detailed information on the UK population, age structure, geography, and more up to date information on the effects of the virus than was available to Report 9. What they found was that CovidSim could closely replicate Report 9, which itself gave an accurate prediction of the first wave. The best match was obtained by using R0=3.5, rather higher than Report 9’s original R=2.4 assumption.

Faced with an infection with a 1% kill rate, “letting it rip” in a population of 66 million produces a death toll of over 600,000 and simply could not be countenanced. Report 9 shows that *at best*, the interventions it modelled could reduce the death toll to around 200,000 over the long run.

For example, the models show that suppressing the virus sufficiently to ensure that ICUs are not overwhelmed (as happened in northern a Italy) does not necessarily reduce the long term death toll if only non-pharmaceutical interventions are available (ie if there is no vaccine or effective therapy). In some cases, interventions can effectively suppress the virus in the short term, but release a larger and more deadly second spike once the interventions are eased, increasing the overall death toll in the longer run.

Report 9 did not model the complete lockdown which was adopted, but was able to model a series of interventions, some of which appear to have bizarrely counter-intuitive effects. For example closing schools or social distancing could in certain (but not all) circumstances increase total deaths. CovidSim replicates these and explains why this can happen.

The UK government’s apparent strategy

The best scenarios in the model result eventually in about 4 times the 60,000 excess deaths seen so far. It predicts a series of waves resulting from lockdowns in response to infections threatening to overwhelm the health system, followed by releases which allow the virus to retransmit. The rapid first wave and now the second slower wave are both well described in the model.

As a consequence, UK policy has been to attempt keep the number of infections down to a level which does not overwhelm the health service and to keep the value of R to just below 1 to keep it that way. It appears to be an attempted compromise between wanting as few casualties as possible and wanting as little restriction on the economy as possible.

Israel, initially successful at controlling its first wave, provides a salutary warning of how difficult in practice that can be. In the face of a bigger second wave, it is now having to impose even stricter measures than before.

From Finding to Isolating

The model shows what happens if you use various interventions as control variables. It did not explicitly include the use of track and trace. At that point, the UK's track and trace capability had already been overwhelmed and was not available as a policy option anyway. It does, however, include "case Isolation" and "household quarantine", which arguably are analogous. A different simulation model developed at Cambridge University led Karl Friston comes to altogether different conclusions about using an effective Find-Test-Trace-Isolate-Support strategy. In this model, relatively modest improvements in FTTIS performance make significant improvements in epidemic control, echoing what has been experienced in practice in New Zealand, Vietnam, Greece, Singapore to name but a few. These countries reacted before the epidemic had taken hold, and it remains to be seen if the UK would, like Wuhan, need to first establish conditions by a total lockdown in which FTTIS became feasible.

So what would be the effect of including a successful find-test-trace-isolate-support system as an available policy variable? Is the current practice of action only on receipt of a positive RT-PCR response the best available? To what extent would we still be reliant on lockdowns as the dominant policy response? Would we still have the principal aim of minimising harm to the economy by tolerating a level of infection which the health service can ostensibly cope with, and be using R=1 as the key control variable?

The meaning of R

To answer those questions, we first need to understand what the value of R tells us, and more importantly what it does not. R is simply the average number of people that someone who is infectious actually infects.

Imagine a country, say New Zealand, which has no cases, but then one turns up and he infects another 6 people before you notice. What is the value of R? Obviously it's 6. If you are quick off the mark, you can isolate all 7 of them and eliminate the epidemic.

Now imagine a country, say UK, which experiences 20,000 new cases in one day, and they go on to infect 30,000 more. What is R in this case? It's 1.5 But finding and isolating all 50,000 before they spread it further is clearly a bit of an ask.

Which do you suppose is more serious in these examples? R=6 or R=1.5? Without a context R doesn't tell you that, only if it's growing or not.

It doesn't tell you how fast it's growing either, although for the same epidemic higher does mean faster. Epidemics, unconstrained, grow exponentially, which means simply that they double in size over a given period. One that doubles every month is clearly less serious than one which doubles every 4 days, because you will have more time to react. R doesn't actually tell you the doubling time directly, though changing R for a given epidemic through policy actions will change it.

Let's imagine a disease A in which on average, an infected person passes it to 4 others, so R=4. Let's suppose though, that a person is circulating and infectious with this for 28 days, so on average he infects someone new once every 7 days. The disease passes on to 2 more people on average in 14 days, so the doubling time is two weeks.

Compare that with disease B in which on average, an infected person also passes it on to 4 others, but over only 8 days, after which he ceases to be infectious. In that case, 2 more people are infected on average after 4 days, so the doubling time is 4 days. Same R, less than half the doubling time.

Two and only two variables to control

In these examples, we have actually been considering the variable R0 (R-nought). It's the value of R for a population that doesn't have natural immunity and doesn't adopt measures to control it. R depends on **both** the probability that you pass on an infection in *any one day*, **and** that you are circulating and infectious for a given *number of days*. R (more formally Rt the rate at time t) is the product of those two controllable numbers (and two more numbers we can't control - see footnote).

So there are just two ways of controlling this epidemic before a vaccine arrives.

Reduce the daily probability of an infectious encounter - lockdowns & behaviour change

The first policy measure is to reduce the probability that on any one day you can pass on an infection. That's precisely the purpose of lockdowns, mask wearing, hand-washing, social distancing, & closing venues where large numbers of people congregate. It's particularly important for a disease that has an asymptomatic infectious phase, because we may not know we are infected even if we would instantly self isolate if we did.

Our experience with Covid-19 in the UK is that a hard lockdown without an effective test and trace reduces R from 3.4 to fractionally below 1, sufficient for prevalence to slowly decrease. But releasing lockdown - opening pubs, restaurants, schools, universities, encouraging people to go back to their offices - spreads the infection exponentially once more, but because people are commonly wearing masks, washing their hands and socially distancing, the doubling time is now somewhere between 8 days and 2 weeks instead of 3 days, with R between 1.3 and 1.6

Reduce the days with encounters - find and isolate the infectious

The second available policy measure is to reduce the number of days for which people are circulating and infectious by quickly finding those people and isolating them. How effective is that likely to be? Let's suppose *for the purposes of simple arithmetic and illustration* that, without any other control measures, Covid has very approximately an R0 of 4 and an infectious period of around 10 days.

Let's also suppose - as the only policy measure - we are able to find everyone who has the infection, and get them all to self-isolate after they have been circulating unawares for 5 days. Now, they have infected only 2 people on average before they are isolated, so R would fall from 4 to 2. Let's suppose we can speed this up, and get to them after just 2 and a half days. Now R falls to 1, and the epidemic won't grow. Speed it up further to find and isolate them all in 2 days, R falls to 0.8 and the epidemic burns itself out.

The effectiveness of a test and trace policy is critically dependent on its speed. Any delays in getting a test appointment, getting the test result, alerting the contact chain, and getting everyone to self isolate will reduce that effectiveness. Using these supposed numbers, take longer than 2 and a half days, and you can't get R below 1 using this measure alone. Right now, RT-PCR test results are taking 2 days to get back from a test site, and 3 days from home test kits. Add in a day to realise that you have ominous symptoms, a day (if you are lucky) to book the test and take it, another day for contact tracers to ask for your list of contacts, and you can see that action to eliminate the chain is unlikely to begin in much less than 5 days. That would correspond to reducing R from 4 to 2 at best in the absence of other measures, which is clearly nothing like fast enough.

So provided FTTIS is fast enough and thorough enough, outbreaks can be contained and eliminated because it is theoretically possible to control R to below 1 so that infections burn out.

Testing, testing

The focus on the *number* of tests is misplaced. The slower the test and trace system, the wider infection spreads and the more tests are needed. The more tests which are needed, the slower test results are likely to be returned as the testing system becomes overloaded. For test and trace to be an effective policy tool, the rate of new infections needs to be within the scope of the testing system to manage. Again, context is everything.

And, of course, *there is no point in any amount of testing if it does not result in the isolation of those infected*, and the isolation of their chain of recent contacts who may be harbouring the virus as yet asymptomatically. It is only recently that the government has introduced any financial support for people “doing the right thing” at an extremely modest level compared to the swingeing fines for anyone who gets a positive test and fails to isolate. Fairly obviously, there will now be a reluctance to take a test for fear of unwelcome consequences due to necessity or unexpected events some time later.

Things can get better

Yet the picture is not completely bleak. Modelling studies indicate that a relatively modest improvement in performance would have a significant effect on further slowing the spread. Where local authority health teams are responsible for contact tracing they have a very good record for finding well over 90% of the contact chain and fairly quickly persuade most of them to isolate. In contrast, the Tier 2 test results managed by SERCO have a much poorer record on all counts. They do pass the “difficult” cases back to the local health teams, but this is only after they have spent time in failing. Local authorities are pressing very hard to take responsibility for all contact tracing in their area, and to be first resort, not last resort, but there appears to be considerable reluctance from within central government to allow this degree of delegation, which many argue would transform performance to the levels seen in other countries where this is the norm.

It’s not R it’s k

Note too that, epidemics happen in small localised clusters and spread from there. So a value of R for the UK is an average number disguising rapidly growing epidemics in some areas in a calm sea of tranquillity in others. So understanding **where** outbreaks are is vitally important, preferably down to the address of, say, the nightclub which is the source of it, or the street in which a number of cases are to be found.

The most significant variable in the epidemiological models may not be R at all, but the measure of dispersion k. In “dispersed” infection, a few super-spreading events/locations/people spread the majority of infections, with a majority of people not spreading it much, if at all. In Covid-19 it seems 80% of infections are spread by just 20% of people, usually at events or locations which facilitate the spread of the virus. This argues strongly that *backward* tracing to find *clusters*, and then immediately isolating all of those involved, irrespective of testing, would be a much more effective public health response than the focus on a centralised system *forward* tracing to find infections one at a time. This is only feasible with local knowledge by local teams. **It’s the main reason why delegation of responsibility to local health authorities is a vital necessity**. That also means that full details of all test results including identity and address of positively diagnosed individuals must be shared immediately with the local health teams, and tracing resources redeployed immediately to fall under their line management control.

Symptoms or tests?

It also questions the reliance on centrally provided RT-PCR tests as the principal means of detecting cases and triggering interventions. Actions triggered by symptomatic diagnosis alone would be far quicker, though that would involve a degree of “noise” caused by misidentification of say flu for Covid-19 until the case could be shown with some certainty not to be Covid-19. There is now evidence, for example, that CT chest scans are showing “powdered glass” lungs highly specific to Covid-19 without RT-PCR registering positive until some while later.

Traditional infection control relied on symptomatic diagnosis and rapidity of response to that alone. The novel factor in Covid-19 is the extent of asymptomatic transmission. To the extent it remains asymptomatic, that’s good news for the one who has it, but bad news for everyone else if he transmits it. Our “accounting” system needs better precision in this. “Cases” (which drive hospital admissions) should be based on symptoms. “Infections” (which drive the epidemic) should be based on positive identification of presence of the virus. We could, for example, get more precise about RT-PCR positive results to winnow out detection of inactive virus particles by following up RT-PCR positives with antibody tests looking for IgG positive, which is indicative that the infectious stage has passed.

Immunity revisited

Now that we have some experience suggesting that re-infection within 6 months appears comparatively rare, and in London, for example 20% may have had the virus, we should revisit the possibility of using volunteers with positive antibody tests within the previous 6 months in critical roles within their capabilities, which are more likely to be exposed or expose others to infection. If there's a choice between back office and front line, put those with likely immunity in the front line, and those more at risk in the back office. Longer term we simply don't know yet.

New and faster test technologies

Rapid RT-PCR testing using LAMP is now available and should be deployed locally (hospital labs, GP surgeries, pharmacies, ports and airports) to reduce the dependency on the slow centralised Lighthouse labs, allowing them to focus on population level surveys and support for local teams without sufficient test capability of their own. Local labs could develop pooled sample RT-PCR testing more easily than established high volume labs, quickly increasing throughput and with the essential rapid turnaround which is required.

Cheap and rapid antigen tests should quickly be brought into widespread use as preliminary screening, even though they are less sensitive and accurate. Positive screening results, or negative results accompanied by suspicious symptoms, should trigger immediate tracing and isolation, and passed through to RT-PCR tests for confirmation. Tracing & isolation can then be stood down if the result is negative.

Putting FTTIS within reach

To reach a low enough level of infection to make a test and trace strategy effective, say 1,000 new infections per day, tight restrictions may be needed temporarily. The tighter the restrictions the shorter they will be required. But responsibility and resources must be transferred away from central control to local authorities where they have a chance of being effective.

This is, in many ways, the disappointment of the summer. Infection rates were falling, and would have halved to be within scope of the testing system within another 2 weeks, and could well have brought prevalence in the North down to that experienced in the South. But then the lockdown was eased to bring R back up to 1, and the testing system now needed to be doubled in scale to cope with the still too high level of infection. It never quite managed to catch up before a further release was encouraged through exhortations to return to the office, Eat Out to Help Out encouraging people to frequent pubs and restaurants, and the re-opening of schools. University term has now tipped the epidemic well into territory requiring a severe lockdown to bring infections under control. It remains to be seen if announced government actions are yet sufficient.

Denys Bennett
14 Oct 2020

Footnote:
R is actually the product of 4 numbers, but this briefing considers are only the two we can alter without a vaccine. The other two are infectiousness and susceptibility:
1 The infectiousness of the disease (eg how many virus particles it takes to infect you) is not alterable, it's a feature of Covid-19.
2 The proportion of susceptible people in the population is also a factor, and this depends on the progress of the epidemic. A change in this value sufficient to make a material difference to R would only come about through failure to control the epidemic with possibly hundreds of thousands of deaths until herd immunity might assert itself, assuming that immunity is not merely temporary. In practice it is this variable that is alterable by mass vaccination, and is the long term means by which it is hoped the virus will be defeated.

References:
Reproducing the long term predictions from Imperial College CovidSim Report 9
https://www.medrxiv.org/content/10.1101/2020.06.18.20135004v1.full.pdf

https://www.theatlantic.com/health/archive/2020/09/k-overlooked-variable-driving-pandemic/616548/