



ILLUSTRATION BY DAVID PARKINS

THE LIMITS OF R

What the reproduction number can and can't tell us about managing COVID-19. **By David Adam**

Mathematicians and public-health experts watched through their fingers in May as British Prime Minister Boris Johnson unveiled a series of charts to explain how the government would guide Britain out of coronavirus lockdown. Perhaps most prominent was a colourful dial with a needle hovering near a single digit: 1.

The dial indicated R , a now-totemic figure in the COVID-19 pandemic. The nation, said Johnson, would set a COVID-19 alert level, to be "primarily determined" by the number of

coronavirus cases, and by R , the reproduction number.

To infectious-disease experts, Johnson's focus on the reproduction number as a guiding light for policy was worryingly myopic. They worry about placing too much weight on R , the average number of people each person with a disease goes on to infect.

In this pandemic, R has leapt from the pages of academic journals into regular discussions by politicians and newspapers, framed as a number that will shape everyone's lives. As Germany's chancellor, Angela Merkel, explained in a widely viewed video this April,

an R above one means an outbreak is growing, and below one means that it is shrinking. In many countries, it is publicly reported every week. In June, epidemiologists at the Harvard T.H. Chan School of Public Health in Boston, Massachusetts, announced a website where anyone can look up the value for any country – and for many smaller regions – in the world.

But fascination might have turned into unhealthy political and media fixation, say disease experts. R is an imprecise estimate that rests on assumptions, says Jeremy Rossman, a virologist at the University of Kent, UK. It doesn't capture the current status of an

epidemic and can spike up and down when case numbers are low. It is also an average for a population and therefore can hide local variation. Too much attention to it could obscure the importance of other measures, such as trends in numbers of new infections, deaths and hospital admissions, and cohort surveys to see how many people in a population currently have the disease, or have already had it.

“Epidemiologists are quite keen on downplaying R , but the politicians seem to have embraced it with enthusiasm,” says Mark Woolhouse, an infectious-diseases expert at the University of Edinburgh in the United Kingdom, who is a member of a modelling group that advises the British government on the pandemic. “We’re concerned that we’ve created a monster. R does not tell us what we need to know to manage this.”

Many policymakers understand this: no one else has linked R so tightly and explicitly to public policy as Johnson did, Rossman says. And despite the coloured-dial chart, it’s not clear how much R is actually driving UK policy. In the weeks after Johnson’s announcement, the government didn’t reference R when it took measures to ease restrictions or lowered the national alert level. (It did not respond to requests for comment for this article.)

But researchers remain concerned that R is looming too large, and is being used for purposes for which it was never intended. “It’s not yet clear what actions they are or are not taking on the back of R . But we are concerned because they’re giving it such prominence,” says Woolhouse.

The origins of R

First used almost a century ago in demography, R originally measured the reproduction of people – whether a population was growing or not. In epidemiology, the same principle applies, but it measures the spread of infection in a population. If R is two, two infected people will, on average, infect four others, who will infect eight others, and so on. The measure allows modellers to work out the extent of the spread, but not the speed at which the infection grows.

Unless they regularly test an entire country’s population, epidemiologists can’t measure R directly. So it is usually estimated retrospectively: disease modellers look at current and previous numbers of cases and deaths, make some assumptions to find infection numbers that could have explained the trend and then derive R from these.

One variant of R , R_0 , assumes that everybody in a population is susceptible to infection. That usually isn’t true, but might be when a new virus, such as SARS-CoV-2, emerges. At the start of the epidemic, assessing R_0 (and other metrics) was crucial for epidemiologists building models of how the disease might spread. But when politicians and scientists talk about R , they usually mean another variant called R_t

(sometimes called R_e , or ‘effective R ’), which is calculated over time as an outbreak progresses and considers how some people might have gained immunity, perhaps because they have survived infection or been vaccinated.

R_t and R_0 both vary with the social dynamics of a population: even an easily transmitted virus will have trouble spreading in a region where people rarely meet. In January, the COVID-19 R_0 in Wuhan, China, was calculated to be between two and three; after lockdown, estimates put the R_t there at just over one (ref. 1).

A lagging indicator

Working out R_t involves trade-offs and compromise. Confirmed cases and mortality figures can be used to infer the total number of infections, but both come with a significant lag – which scientists estimate could be anything from a week to three weeks or more. “If

“ R does not tell us what we need to know to manage this.”

you have your R_t estimate lagging by at least ten days, possibly two weeks, then it’s not going to be that useful as a real-time decision-making tool,” says Gabriel Leung, a public-health scientist at the University of Hong Kong.

With a mathematical trick called nowcasting, researchers can use the observed statistical distribution of reporting delays to predict how much higher the number of fresh infections will be in, for example, two weeks. Some estimates of R_t already rely on nowcasting infection data in this way: it is “the method with the least guesses”, says Lars Schaade, vice-president of the Robert Koch Institute in Berlin, Germany’s main public-health agency, which reports a daily and seven-day R_t value based on infections reported by state health authorities.

Nowcasting infections on the basis of trends in past COVID-19 cases is tricky enough, but mortality data typically come with a longer lag, because of the extra time someone has the disease before they succumb to it and because of the paperwork involved in registering deaths, which can take weeks or months to file. A group led by Sheila Bird at the University of Cambridge, UK, publishes nowcast data of COVID-19 deaths in English hospitals. But they cannot yet do the same with a separate data set of deaths compiled by the Office of National Statistics (ONS) because the researchers don’t have access to the necessary data on registration delays: the time difference between when a death occurred and when the ONS reported it.

Extra uncertainty

An issue with nowcasting is that it swaps one problem for another, says Sebastian Funk, a disease modeller at the London School of

Hygiene and Tropical Medicine, who is also advising the British government on this pandemic. “You can try to do that, but for obvious reasons it always comes with uncertainty. There’s no way that you can know how many cases would still be observed that have already been infected,” he says.

Other data on the pandemic’s progress can feed into estimates of R_t by serving as proxies for infections and social behaviour. One is hospital and intensive-care admissions. Another is results from random testing of a population to see how many people currently have COVID-19, or have had it. Researchers also conduct contact surveys, which ask people who they mix with, and can be used to infer changes in R on the basis of estimates of how many others an infected person could meet, although these are time-consuming and could cover only small groups of people. Contact surveys in China showed daily contacts were reduced by seven- to eightfold during the COVID-19 social-distancing period, when most interactions were restricted to the household². Another way to observe trends in people’s movements is to use location data based on the signals from mobile phones, published by Facebook and Google.

“There’s a bit of a trade-off here,” says Funk. “There are some methods that are more immediate but not epidemiological, and there are others that are more directly epidemiological but at the same time more out of date.”

Groups of epidemiologists, Funk says, each have their own approach to combining and using these disparate sources of data to work out R_t , relying on their own statistical models to look at trends in presumed infections. To calculate the official R_t of the United Kingdom, about ten groups present the results of their models to a dedicated government committee, which reaches consensus on a possible range. The figures are presented in that range (currently 0.7–0.9), showing how uncertain the estimates are, but the individual models are not released.

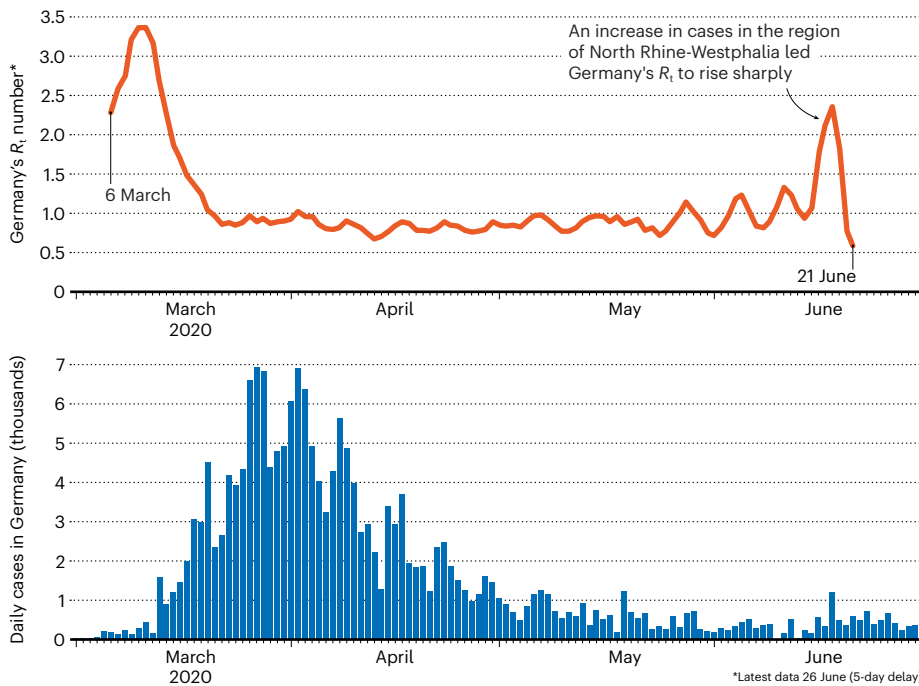
Unofficial estimates

Those ‘official’ R_t numbers are not the only versions available. Academic researchers have taken advantage of infection and mortality figures collated by the World Health Organization and independent groups such as the Coronavirus Resource Centre at Johns Hopkins University in Baltimore, Maryland, to publish R_t figures for numerous countries and states. In late April, for example, public-health researchers in Colombia claimed that the R_t for the first ten days of the pandemic was above two in seven Latin American countries³. The Harvard researchers’ website currently estimates that R_t is above one in more than 30 US states.

Even non-experts can use plug-and-play formulas to create their own variants of R – which can sometimes lead to problems. In May, local newspapers across England ran

GERMANY'S REGIONAL OUTBREAKS

Germany's R_t spiked above two when the country had a low level of COVID-19 cases, largely because of a cluster of local infections in one region.



stories claiming to reveal regional R_t values for specific towns and cities. The *Swindon Advertiser* claimed the town's R_t was 0.35, perhaps "one of the lowest in UK". But officials at Brighton and Hove City Council (labelled with the fourth-highest R_t , at 1.7) issued a statement calling the figures misleading and potentially dangerous. "It is not possible to calculate meaningful R values at a very local level," said Alistair Hill, a public-health official on the council.

The figures were not, it turned out, R_t values at all: they came from an index created by the founders of a London-based analytics start-up called deckzero.com. That index, termed R_z , was intended to show how fast local epidemics were growing on the basis of case data from local authorities; it is not an established variable in epidemiology, says Jenna Wang, a co-founder and director of the firm. On 7 June, the founders withdrew their page from public access and said it had been "interpreted out of the context and scope of its original intention".

The drawbacks of an average

An important aspect of R_t is that it represents only an average across a region. This average can miss regional clusters of infection. Conversely, high incidences of infection among a spatially distinct smaller subsection of a population can sway a larger region's R_t value. For instance, Germany's national R_t value jumped from just over 1 to 2.88 in late June (later revised down to 2.17) largely because of an outbreak in a meat-processing plant at Gütersloh in North Rhine-Westphalia (see 'Germany's Regional Outbreaks'). The Robert Koch Institute noted

that national infections overall were still low, which is why the local outbreak had such an effect on the country's R_t , which had dropped below 1 again by the end of June. This makes it unlikely that R_t would be used to steer local lockdown policy in Germany, Schaade says. "If the rolling mean of R was at 1.2 for a few weeks, then that would show there was a problem that needed attention, even if case numbers were low." But in practice, researchers find out

"It is not possible to calculate meaningful R values at a very local level."

about local outbreaks before that because of a reported spike in cases, not because of changes to R_t . Germany has ongoing surveillance and public reporting of transmission levels in 400 counties.

And most experts say that the R_t for the United Kingdom is kept artificially high by the very large numbers of infections and deaths in care homes for older people, and does not reliably represent the risk to the general population.

Regional R_t numbers have been touted as a way to guide the further easing of restrictions, because they could allow a place that showed a resurgence in cases to be sealed off. But regional R_t numbers become less accurate as they are applied to smaller populations, especially when absolute infections are low.

The Harvard site produces numbers for US counties – which can range from thousands

to millions of inhabitants – but one of its creators, Xihong Lin, says that hyperlocal data come with big uncertainties. The researchers don't calculate an R_t for a county unless there are ten cases, Lin says. And she stresses that policymakers should not use them in isolation, but only alongside other measures such as the total number of cases and whether it is increasing. "When making recommendations, it's definitely important to look at the whole picture and not just rely on R_t ," she says. Used properly, the data could help public-health officials to identify hot spots of infection to prioritize resources such as testing, she says.

No accounting for superspreaders

Another subtlety not captured by R_t is that many people never infect others, but a few 'superspreaders' pass on the disease many more times than average, perhaps because they mingle in crowded, indoor events where the virus spreads more easily – church services, choir practices, nightclubs and birthday parties, for instance. As few as 10–20% of infected people seem to cause 80% of new COVID-19 cases, Leung says. (Epidemiologists describe this using a 'dispersion' parameter, K , which depicts the variation in viral transmission among infected hosts). That means bans on certain crowded indoor activities could have more benefit than blanket restrictions introduced whenever the R_t value hits one.

When countries consider when to reopen schools and offices, a key question is not only R_t , but what the actual number of infected people walking around is. Denmark and the United Kingdom have similar R_t values for instance, but because the number of infected people walking around Denmark is ten times lower, it's safer for their schools to be reopened.

"When infection numbers are low, maybe you don't care so much about what the reproduction number is, or at least don't care if there's some uncertainty in it," says Funk. A test for the United Kingdom, says Woolhouse, will be whether the country overreacts if case numbers are low but modellers estimate that R is above one.

All that demotes the usefulness of R in deciding policy, say Funk and others. For countries recovering from the first wave of the pandemic – such as the United Kingdom – researchers say it's far more important to watch for clusters of cases and to set up comprehensive systems to test people, trace their contacts and isolate those infected, than to watch the needle swinging on a colourful dial.

David Adam is a science journalist based in London.

1. Kucharski, A. J. et al. *Lancet* **20**, 553–558 (2020).
2. Zhang, J. et al. *Science* **368**, 1481–1486 (2020).
3. Caicedo-Ochoa, Y. et al. *Int. J. Infect. Dis.* **95**, 316–318 (2020).

SOURCE: ROBERT KOCH INSTITUTE/JOHNS HOPKINS