5
The Mathematics of COVID-19

Michael Edesess

Emerging infectious diseases represent major threats to public health. SARS-CoV-2, a novel coronavirus that became known as the COVID-19 disease, not only resulted in very high numbers of infections and deaths all around the world, but also led to huge economic losses and social disruptions that threatened global security. While every pandemic is unique—and COVID-19 certainly has a set of unique manifestations as noted in Chapter 1—when a new disease, or a new variant of an old one, bursts upon the scene, four questions urgently need answers:

- How does it spread and how fast will it spread?
- How many cases will require hospitalisation or other emergency medical attention, and when will these needs arise?
- How many deaths will there be?
- What can be done to change and improve the outcome?

Mathematics played an important role in helping policy-makers and healthcare professionals answer these questions when COVID-19 emerged. Mathematical modelling provided quick, approximate answers. Its predictions improved as more information was gathered that could be applied to help stem the course of the disease. Over the course of COVID-19, including its variants, mathematics contributed greatly to fighting the pandemic. Much learned research was produced all around the world, resulting in many publications on all aspects of the disease.

But mathematics helps in another way. When a disease first appears and begins to spread, the situation is like the proverbial 'fog of war'. Little is known and it is confusing to try to understand what really matters and what doesn't. Mathematical modelling requires attention only to those variables that actually affect the spread of the disease. It tends to focus on the things that matter and removes focus from those that do not, providing much-needed clarity. For example, as we will see, from a mathematical modelling standpoint, the only things that matter to the spread of disease are (1) the number of contacts per day between an infectious person and a susceptible person; (2) the probability that the disease will be transmitted during a contact; and (3) the

number of days for which an infectious person is infectious. Focusing on those things, as a mathematical epidemiological model does, helps policy-makers determine what levers they need to push to affect the course of the disease.

While epidemics emerge at least every five years, few of them have caused the high level of global concern, strong action, and urgency that COVID-19 did. The reason is simple: COVID-19's combination of transmissivity and virulence led to high hospitalisations and deaths. Other epidemics were contained quickly, had low transmissivity or low virulence, died out on their own after an initial panic (such as severe acute respiratory syndrome or SARS), submitted to vaccines or pharmaceutical remedies, or became endemic in the global population with varying effectiveness of treatments and preventive measures (such as malaria or tuberculosis). What was unique about COVID-19 was that mathematical models predicted a very high level of severe cases, hospitalisations, and deaths in the tens of millions if immediate action was not taken. These predictions were taken very seriously in some countries and acted upon, and much less seriously in others, partly because in those countries the mitigating measures were not acceptable to many people and to the countries' leaders themselves. This difference in national action in response to the predictions, and the explanations for them, has widened the gulfs between governmental systems and social contracts, and even between social groups within nations. The key catalyst was the predictive models, upon which countries could either act firmly, rapidly, and decisively or in a more desultory manner, depending on their national philosophies and level of organisation.

How Fast Will It Spread?

At the beginning of any outbreak, the speed of spread is exponential. Exponential growth, a concept taught at school, starts slowly at first but then is extremely rapid. Exponential growth is often illustrated by the example of placing coins, or grains of rice, on a checkerboard with sixty-four squares. One coin is placed on the first square, two on the second, four on the third, and eight on the fourth. The number of coins placed on a square then continues to double on each square. Most people are completely surprised to learn that by the time the process gets to the 64th square, the number of coins placed on that square will be more than 18 million trillion, a number that can be written as the number 18 with 18 zeros after it. If they were grains of rice instead of coins, the quantity of rice placed on the last square would weigh more than 387 billion tonnes, about the weight of Mt. Everest.

In the case of the spread of disease, suppose for example that the number of COVID-19 infections doubles every five days. This was roughly the case before any measures were taken to stem the spread. As Figure 5.1 illustrates, what doubling every five days means is that if the disease began with one person, two people would have been infected in five days; four in ten days; sixty-eight people in a (30.5-day) month; 4,700 people in 2 months; 323,000 in 3 months; and 22 million in four months.

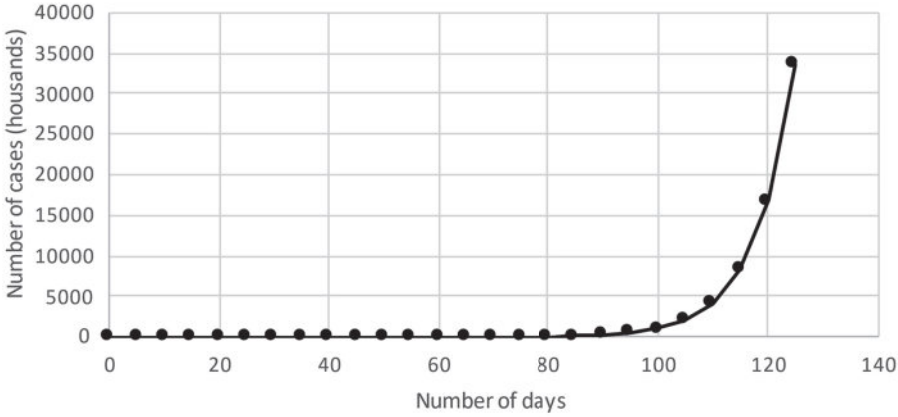

Figure 5.1: Exponential growth in case numbers

Suppose that 10 per cent of cases require hospitalisation. That would mean that at least two million hospital beds would be needed in four months—a tenth of the 22 million cases. If there were only 900,000 hospital beds available (approximately the actual number in the United States in 2021)[1] that would mean more than a million COVID-19 patients would not have hospital beds. Besides hospital beds, doctors and nurses are also essential, but there would be a shortage. There would be other patients that needed hospital beds and professional care for other serious medical problems. Such a situation would truly be a disaster. In addition to the deaths that would occur, what the healthcare community was most concerned about was extremely overburdened medical facilities, supplies, and service professionals. Using China as an example, there were only 6.41 public health professionals per 10,000 population at the end of 2019.[2] In other words, it could lead to the health system being totally overwhelmed.

There are only two solutions to such a situation, which should preferably be deployed concurrently: construct more hospitals and expand the number of health workers very quickly, and constrain the disease from spreading. Since there is a limit to building hospitals that quickly, measures had to be taken to stem the growth. It should be noted that China developed the *fangcang* cabin facilities that could be built rapidly, and healthcare professionals and personal protective equipment were mobilised and dispatched from other parts of the country to send to Wuhan and Hubei during the initial outbreaks (see Chapter 9), which helped the situation, but reducing the growth rate of COVID-19 remained the goal.

1. American Hospital Association, 'Fast Facts on U.S. Hospitals, 2022', accessed 13 September 2022, https://www.aha.org/statistics/fast-facts-us-hospitals.
2. Shuangyi Sun, Zhen Xie, Keting Yu, Bingqian Jiang, Siwei Zheng, and Xiaoting Pan, 'COVID-19 and Healthcare System in China: Challenges and Progression for a Sustainable Future', *Globalization and Health* 17, no. 1 (2021): 1–8.

The Urgency of Reducing the Rate of Growth

If the growth rate of the infectious disease could be cut in half, then in the example above, the number of hospital beds needed in four months would be only 470 (10 per cent of 4,700)—the same as the number required in two months under the fast-growth scenario—instead of two million. This shows what a huge impact reducing the rate of spread by half would make. If the disease continued to spread at the halved rate, there would still be a need for 2 million hospital beds in eight months instead of four. Perhaps by then, a medicinal cure or a vaccine would be found—or the disease might have dissipated on its own by mutation. That would be the hope. Meanwhile, the strain on hospitals, medical care workers, and equipment would have eased.

The Determinants of Spread

The answer to the following question is the key to determining how fast the disease will spread.

If a given person, let's call him Bob, has the disease and is infectious, how many other people will get the disease from Bob? The answer depends on the answers to three further questions:

- How many susceptible people (people who can catch the disease) does Bob interact with each day?
- How likely is it that each susceptible person will catch the disease from Bob when Bob interacts with them?
- For how many days is Bob infectious?

For example, suppose Bob interacts with an average of ten people a day, and they are all susceptible to the disease. Let's assume none of the ten people has had the disease and so have no immunity, and that there is as yet no vaccine or cure available. Suppose that each time Bob interacts with someone, let's say Alice, there is a one in 20 chance that Alice will catch the disease from Bob. That is, the probability that Alice will catch the disease from Bob during their interaction is 0.05. To put it another way, for every 20 susceptible people that Bob interacts with, one catches the disease.

And suppose further, for our example, that Bob is infectious for six days. Thus, Bob interacts with ten people each day for six days. With each interaction, there is a one in 20 chance that the person will get the disease. Multiplying these three numbers together shows how many people Bob will infect: 10 people/day, times 0.05 chance of infection, times 6 days = 3 people ($10 \times 0.05 \times 6 = 3$). The result of multiplying these three numbers together derives what is called R, the *reproduction number*—that is the average number of people each person with a disease goes on to infect. Hence, in this example, one infectious person will infect three more people—that is $R = 3$. If Bob was the *first* person to be infected, so that every person that Bob interacts with is susceptible to the disease, then R is designated $R0$ (pronounced R-zero or R-naught). $R0$ is called the *basic reproduction number*.

It is assumed in this example that once infected, the person will not be susceptible to it for the duration of the simulation period, which is usually no more than a few months. If Bob is not the first person infected, then some people he interacted with would have had the disease already and would not be susceptible any longer. When that happens, *R* becomes less than *R0*. More about that later.

R can also become less than *R0* because measures are undertaken to stem the spread. Public policy is focused intensively on reducing *R*.

Measures to Stem the Spread of the Disease

Assuming no treatment or cure is yet available, the number of days for which the diseased person, Bob, is infectious cannot be changed. That leaves only two variables that can be altered to reduce the number of people Bob infects:

(1) the number of people Bob interacts with each day; and
(2) the probability that Bob gives the disease to a person when Bob interacts with them.

The number of people Bob interacts with each day can be reduced by isolating or quarantining Bob to keep him away from other people. The probability that Bob gives the disease to another person when he interacts with them can be reduced if Bob wears a mask and keeps a distance from the other person of at least one and a half metres.

Neither of these is an absolute guarantee, of course. A very strict quarantine, however—such as was adopted in mainland China—is almost an absolute guarantee that Bob would not interact with anybody while he is infectious, and therefore would not give them the disease.

Less strict isolation or quarantine policies, such as were practised in many countries during the COVID-19 pandemic, provide less of a guarantee that an infected person would not interact with other people, but they did reduce the number of interactions enough to have an important effect. Thus, if Bob can be induced or required to interact with only five people a day instead of ten, that will cut the growth rate in half and hugely slow the rate of spread.

In addition to the possibility of infection due to direct in-person interactions with infected persons, there is the possibility of catching it from touching viral residues on surfaces and then touching one's mouth, nose, or eyes. The risk of infection is reduced by sanitising surfaces and washing hands. The route of transmission via surfaces was regarded as important during the first six months of the COVID-19 outbreak but it was determined subsequently that the chances of contracting the coronavirus from surfaces were low, and the major transmission route by far was in-person interactions.[3]

3. U.S. Centers for Disease Control and Prevention, 'Science Brief: SARS-CoV-2 and Surface (Fomite) Transmission for Indoor Community Environments', https://www.cdc.gov/coronavirus/2019-ncov/more/science-and-research/surface-transmission.html.

Using the Reproduction Number to Predict the Spread

In the simplest possible form of the standard mathematical model, everyone is assumed to have the same *R*—that is each infected person is assumed to infect the same number of other people. An important additional number is the *generation time*. Generation time is the time from when a person gets infected until the next person that person infects becomes infected. The generation time for SARS and COVID-19 has been estimated at seven days. Hence, the first person infected, Bob, will infect three additional people in seven days. Each of those people will infect three more people in another seven days, for a total of nine people infected after 14 days (in addition to the original three). And after 21 days, 27 more people (3×9) will be infected.

Figure 5.2 shows the growth of infections over the course of four months if the population were always 100 per cent susceptible (or if the population were infinite).

Of course, the population is not infinite in real life, and will not always be 100 per cent susceptible either. The percentage of the population that is susceptible will decline over time, as people become infected and recover with immunity, or die as a result of being infected. When an increasing percentage of the population is no longer susceptible, the reproduction rate *R* declines because the number of susceptible people Bob or another infected person interacts with is less. This causes the number of infections to eventually peak, and then decline, as shown in Figure 5.3. Figure 5.3 assumes the total population is 330 million, like the population of the United States, and it is assumed that there are no cases brought in from outside the country.

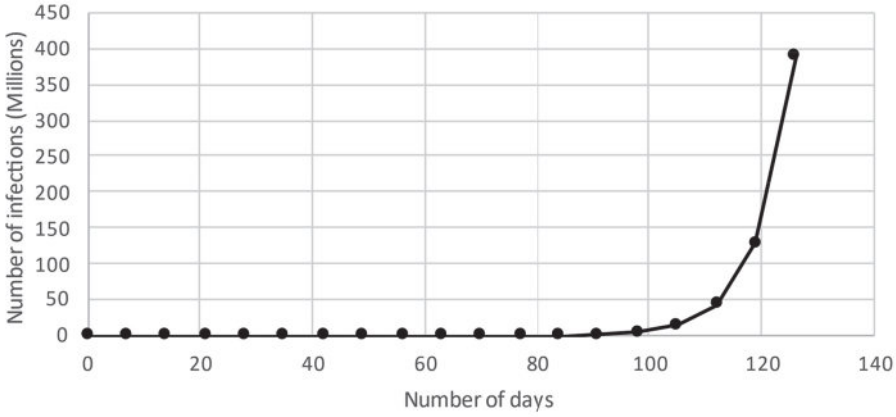

Figure 5.2: Growth of infections, infinite population

In Figure 5.3, the dashed line is the number of cases added each day, while the solid line is the number of active cases—that is, those that are still viral. If the hospitalisation rate is 10 per cent, then the number of people with the disease in hospital would be 10 per cent of the values on the solid line.

In Figure 5.4 the assumed population is 7.9 billion, which is the world's population today. With a larger population, it takes longer for the disease to peak—about four months in this example instead of three and a half.

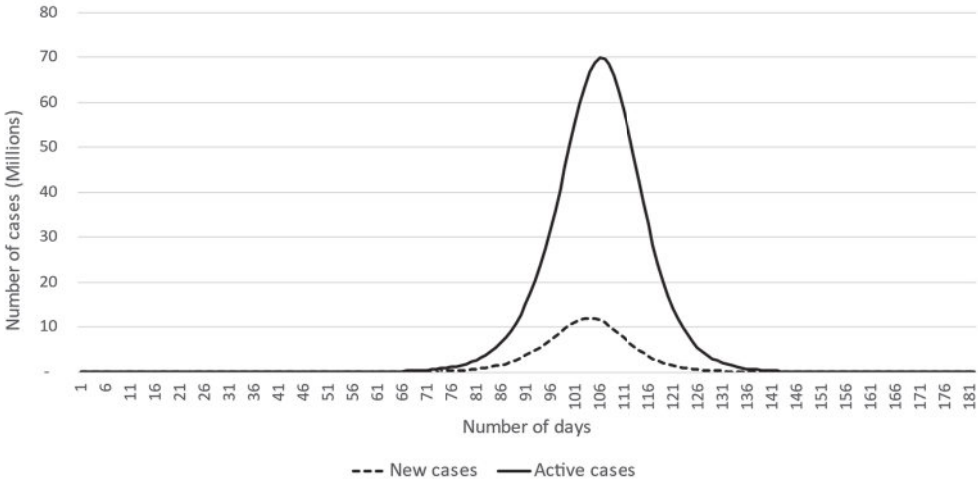

Figure 5.3: Predicted cases over time—United States size population

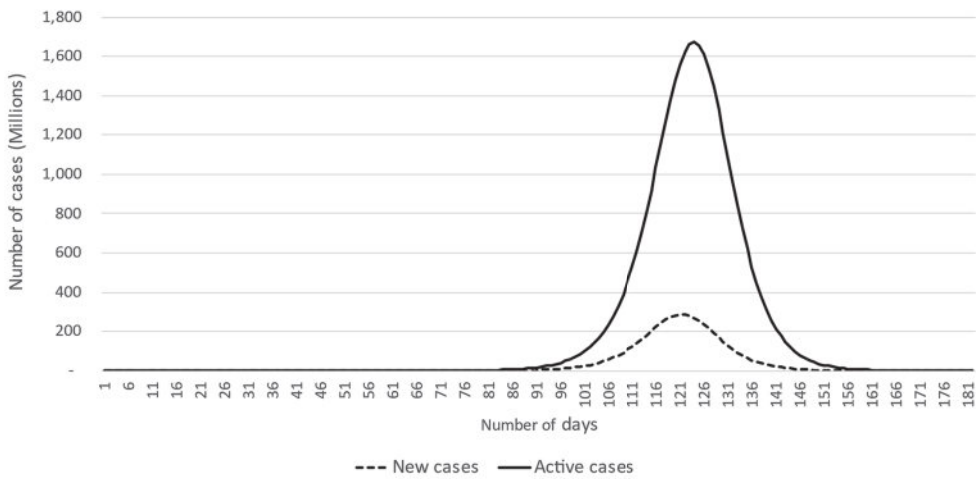

Figure 5.4: Predicted cases over time among world population (7.9 billion)

In Figure 5.5, the cumulative number of cases is also shown, as a dotted line. Figure 5.5 shows that in this model, if nothing were done to reduce the growth rate of the disease, ultimately almost six billion people would have contracted it, about 75 per cent of the world's population.

Figure 5.6 shows why the number of cases peaks and then declines. In addition to the number of daily new cases and currently viral cases over time, it also shows the *R* number, with its value on the right axis.

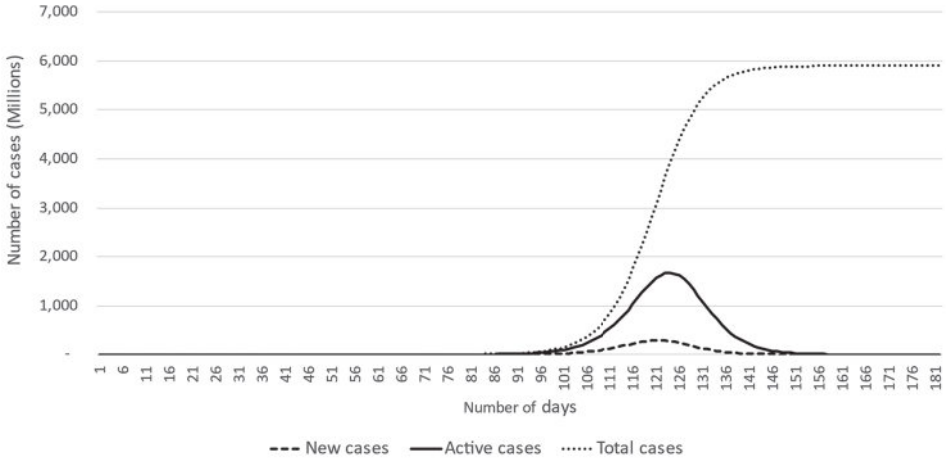

Figure 5.5: Predicted cases over time among world population (7.9 billion)

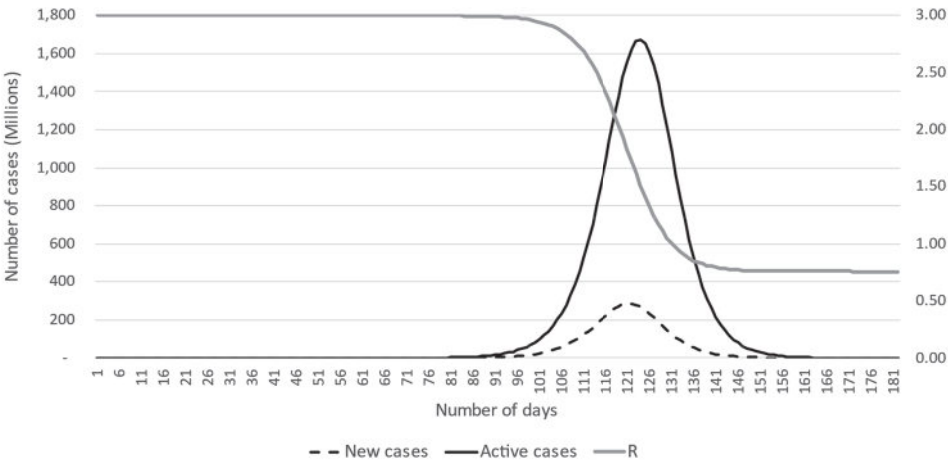

Figure 5.6: Predicted cases over time among world population (7.9 billion)

Notice how *R* declines as the disease catches fire in the population. At the beginning and for a long while, before exponential spread really takes off and the disease afflicts a large number of people, *R* has a value equal to or very close to its initial value of *R0* = 3 (each infectious person infects three more people). As a larger and larger percentage of the population is infected and is no longer susceptible to infection, *R* declines because the number of susceptible people an infected person interacts with is now less.

Finally, as the rightmost part of Figure 5.5 shows, the disease no longer grows in the community, even though (in this example) 25 per cent of the population is still susceptible. This is known as *herd immunity* (see Chapter 4). What happens is that, as the number of people infected starts to decline, the percentage of the people they meet who are susceptible declines too. Eventually, they are unable to meet each other before the infected people are no longer infectious anymore, and the disease dies out.

The SIR Model

The graphs above are an example of the kind of output produced by the Susceptible-Infected-Removed (SIR) model. This is the standard model used by most mathematical modellers of the spread of disease. In this model, the population begins with all but one person, the first one infected, susceptible to the disease. Then, as the first infected person and gradually many infected people start to interact with the susceptible people there is a chance that a susceptible person will become infected. This chance is called the *transition rate* from susceptible to infected. It is often measured—or estimated—as a daily probability. Once infected, an infected person has a daily probability of transitioning to 'removed', which can mean either recovered or dead. In either case, that person is no longer in the susceptible pool.

The examples shown above are only illustrative and do not represent actual predictions that were made by any specific SIR models. But all predictive models will have the same pattern over time, at least in the absence of public policies to alter the predictions.

But All *R*s Are Not the Same

Most SIR models do not make the simplistic assumption that every infected person has the same *R*. 'Compartmental' models place people into different compartments, at different times of day, in different locales and engaged in different kinds of activities, where they will have different *R*s when they interact with other people within the compartment, and yet other different *R*s when they interact with people in other compartments.

For example, in one model of the spread of influenza, people were put into these different compartments: child in household; adult in household; child in small play group; child in large day-care centre; child in elementary school; child in middle school; child in high school; adult in workgroup; adult in neighbourhood; and adult

in community.[4] Different contact probabilities were assumed for each pair of possible contacts. For example, the probability of contact (per day) of two children in a household was 0.6 (60 per cent chance of contact), while the probability of contact of two children in a large day-care centre was only 0.15, or one-fourth as much. (A contact was defined as being within a specified distance of each other for a specified length of time.) Different probabilities of contact yield different *R*s. The model also needs to make assumptions about how much time a child, for example, spends at the day-care centre (or school), how much time in the household, and how much in the neighbourhood or community. Many assumptions are needed to be fed into a full-scale, advanced SIR model. The assumptions are, of course, of necessity imprecise, but they are the best that can be made.

Running the SIR Model Base Case

Once all these assumptions are fed into the model, it can be run for a large population that is allocated to the various compartments (another set of assumptions, usually obtained from demographic data). Running the model entails beginning with one or only a few infected cases, then simulating the progress of the spread day by day after that. Each day, some proportion of the susceptible people will transition to infected, and then, some proportion of the infected people will transition to the removed category. This will provide how many of the population are still susceptible, infected, or removed on each future day.

The model will also make additional assumptions about how many of the infected will be hospitalised, and how many will die. These assumptions may be different for different age groups. Hence, the models can make a prediction not only about how many people will be infected on each day in the future, but how many will be hospitalised and how many deaths there will be.

Hypothesising Public Health Policies and Changing the Assumptions Accordingly

The mathematical modelling base case is run under the assumption that nobody changes the way they lived their lives before the disease started to circulate. This is of course not a realistic assumption, but it is standard practice for modelling. In reality, people would likely change their routines and their number of interactions with other people out of fear of the disease. However, that is not likely to reduce the spread enough. Up until March 2020, in the United States and the United Kingdom, the public policy approach was to do practically nothing.

4. Timothy C. Germann, Kai Kadau, Ira M. Longini, Jr., and Catherine A. Macken, 'Mitigation Strategies for Pandemic Influenza in the United States', *Proceedings of the National Academy of Sciences* 103, no. 15 (2006): 5935–5940 (supplemental materials).

This changed dramatically after the results of research by academics at Imperial College, London, were announced in mid-March 2020.[5] That research predicted 510,000 deaths in the United Kingdom and 2.2 million in the United States if nothing were done to mitigate the spread of COVID-19. That caused concern. The Imperial College study also explored how those numbers of deaths could be reduced if certain public policy interventions were adopted to contain the spread. Those interventions included: case isolation in the home; voluntary home quarantine; social distancing of those over 70 years of age; social distancing of the entire population; and closure of schools and universities.

Compared to measures that had already been taken in China beginning in late January, these were mild measures. In China, much stricter measures were taken to try to ensure that the disease did not spread any further at all, after the first few weeks of spread. These measures were notably successful, as discussed in Chapter 9.

In the Imperial College study, further assumptions had to be made about how much each of the potential policies to contain the spread of the disease would reduce the *R*s. Then for each possible containment policy, and combination of policies, the Imperial College team reran the model. Figure 5.7 shows the predictions from the Imperial College study's results for several different policy measures and combinations thereof.

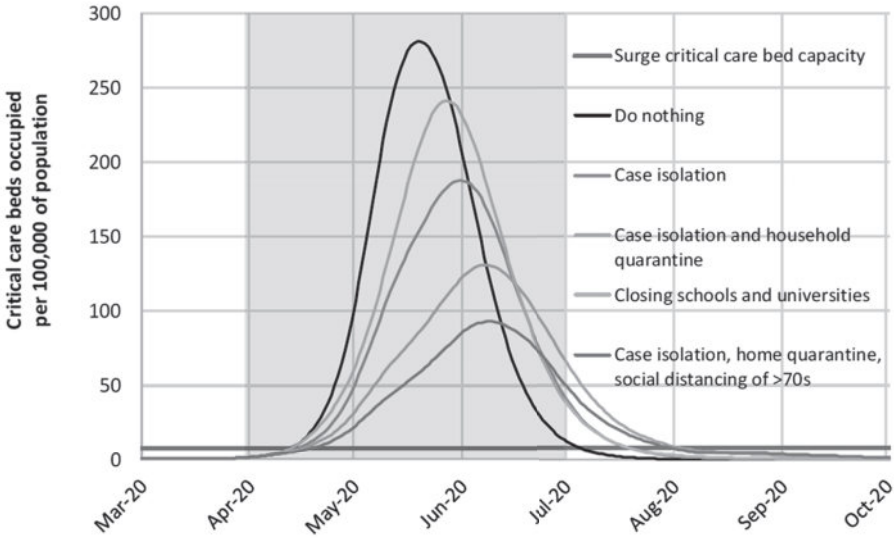

Figure 5.7: Mitigation strategy scenarios for the United Kingdom showing critical care bed requirements. Courtesy of Neil M. Ferguson et al.

5. Neil M. Ferguson, Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, et al., 'Report 9: Impact of Non-pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand', Imperial College COVID-19 Response Team, 16 March 2020.

The solid black ('Do nothing') line (line with the highest peak number of critical care beds) shows the unmitigated epidemic (base case). The light grey line (line with the second-highest peak) shows a mitigation strategy incorporating the closure of schools and universities; the line with the third-highest peak shows case isolation; the line with the fourth-highest peak shows case isolation and household quarantine; and the line with the lowest peak shows case isolation, home quarantine, and social distancing of those aged over 70. The shaded area indicates the 3-month period in which these interventions are assumed to remain in place.

In order to arrive at these predictions, the modellers had to make assumptions about how much each of the interventions would reduce the rate of spread. For example, they assumed that for 'case isolation in the home', symptomatic cases would stay at home for seven days and that this would reduce non-household contacts by 75 per cent during that period. They also assumed that 70 per cent of households would comply. For 'social distancing', they assumed that it would reduce contact rates by 50 per cent in workplaces and reduce other contacts by 75 per cent, but that as a result, it would increase household contacts by 25 per cent (because people would be at home more), and they assumed 75 per cent compliance with the policy.

Each of these assumptions for a mitigation policy changes the contact probability assumptions and the *R*s when the model is run. This is how the alternative sets of predictions for different mitigation strategies are arrived at in the modelling process.

Interventions to Reduce the Rate of Spread

As mentioned before, there are two ways to reduce the rate of spread: reduce the number of contacts an infected person has and reduce the probability that the person contacted will catch the disease. Reducing the probability that a person will catch the disease from an infected person is relatively straightforward—wear a mask (and possibly other protective gear) and maintain a distance. Therefore, almost all of the intervention strategies have one objective: to reduce the number of contacts made between infected persons and susceptible persons.

The first priority is to identify infected individuals. This can be done by means of testing for COVID-19 and tracing the contacts of anyone who tests positive. Beyond that, it is all about isolating and quarantining anyone who either has tested positive for the virus or has been in contact with someone who tested positive, or even someone who was in contact with someone who was in contact with someone who tested positive—unless they have repeatedly tested negative. How effective these strategies are, depends on how strictly they are enforced, or adhered to.

Superspreaders

As noted above, not every infected person has the same *R*. As a matter of fact, studies have shown that the dispersion of *R*s among infected individuals is very wide. This

dispersion is measured by another letter of the alphabet, *k*. (Confusingly, a small *k* indicates wide dispersion—it has been estimated that *k* for COVID-19 has a low value of 0.1).[6] It has been found that some infected people, and some gatherings of people, contribute to the *overdispersion* of *R*s. In other words, there seems to be a small percentage of the infected who have very high *R*s, whether because they carry a high viral load or because they interact at close quarters with a large number of people. It appears that a small percentage of infected people do most of the spreading of the virus, while the much larger percentage spread it relatively little. For example, a two-and-a-half-hour chorus practice in the American state of Washington in May 2020, attended by sixty-one persons among whom there was one person infected with COVID-19, resulted in at least thirty-two additional cases and perhaps as many as 52, when secondary infections are considered.[7]

Both people and events that spread the virus unusually widely are referred to as *superspreaders*. The importance of the phenomenon of superspreading—both superspreading individuals and superspreading events—is that it has implications for contact tracing.

Contact tracing has typically been done when a person is confirmed to be infected. They are then questioned as to which other people they interacted with and what venues they have been to since they got infected. This way, people who may have caught the virus from them can be identified, tested, and isolated if infected. This is called *forward tracing*, because it identifies contacts going forward in time beginning with when the person became infectious.

The fact of superspreading events and people indicates that more cases can be winnowed out by doing *backward tracing*. This means that in addition to identifying with whom the infected person has interacted *since* becoming infectious, the investigation goes back to the event or person from whom the infected person contracted the disease. Because the infecting person or event may have been a superspreader, the backward tracing process seeks to identify who else may have contracted the virus from the superspreader.

Figure 5.8 shows why more cases of infection are discovered by doing both forward and backward contact tracing.[8] Black dots indicate detected cases, dark grey dots quarantined cases, and light grey dots undetected cases. This chart shows two infectious cases are discovered, 'Index case #1' and 'Index case #2' (dark grey dots to the left and right of chart A). They have a common source in a 'Primary case', but that primary case

6. Akira Endo, Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group, Sam Abbott, Adam J. Kucharski, and Sebastian Funk, 'Estimating the Overdispersion in COVID-19 Transmission Using Outbreak Sizes Outside China', *Wellcome Open Research* 5 (2020).
7. Lea Hamner, 'High SARS-CoV-2 Attack Rate Following Exposure at a Choir Practice—Skagit County, Washington, March 2020', *Morbidity and Mortality Weekly Report* 69 (2020).
8. See Akira Endo, Quentin J. Leclerc, Gwenan M. Knight, Graham F. Medley, Katherine E. Atkins, Sebastian Fun, et al., 'Implication of Backward Contact Tracing in the Presence of Overdispersed Transmission in COVID-19 Outbreaks', *Wellcome Open Research* 5, no. 239 (2020), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7610176.3.

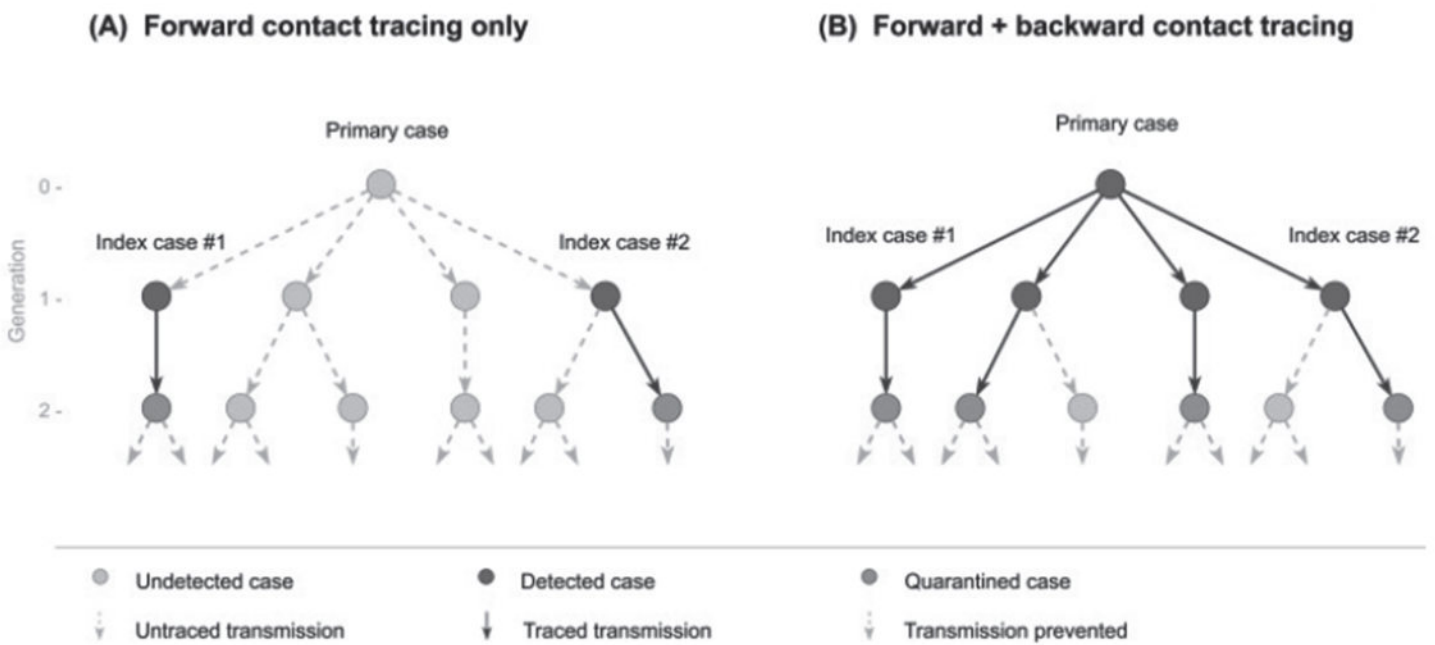

Figure 5.8: Forward and backward contact tracing. Courtesy of Akira Endo.

is initially undetected (light grey dot at the top of chart A). In (A), only forward tracing is conducted, identifying only two infected persons, while those coloured light grey are undetected. In (B), backward tracing is conducted in addition to forward tracing. Backward tracing identifies the primary case, which therefore has become a black dot, and then forward tracing from there identifies additional contacts that were made with potentially infectious people (there are still two light grey undetected cases in (B) because contact tracing is imperfect). Some of those additional contacts tested positive and are quarantined.

The Actual Pattern of Cases and Infections over Time

The results of simulations shown in Figures 5.3 through 5.7 do not represent what happens in the real world. They show only what will happen if a single course of action—or no action—is pursued without deviation. In the real world, actions taken in response to the disease change over time, and the disease can change too.

The best analogue to the spread of a disease is the spread of a wildfire. If the wildfire finds a patch of kindling or dry wood or dry shrubbery it can spread extremely quickly. If that fire is then put out but not completely extinguished, so that it smoulders for a while afterwards, its smouldering remains can again find or leap to another patch of kindling or dry wood or shrubbery. A fire that is not completely extinguished can even smoulder underground, undetected, and emerge at a distance to flare up very rapidly again. The spread of a viral disease is similar if it is not completely extinguished. If there are remaining viruses lurking in the population then it can flare up again, astonishingly quickly, just as it could at the onset of the disease.

China is the most prominent exception among countries. For other countries, the objective of their intervention measures was not to eradicate COVID-19 completely, but to 'flatten the curve', meaning to reduce the level of the predicted peak of cases, hospitalisations, and deaths. A key objective was to get the number of hospitalisations and the demand for intensive care units, ventilators, and other specialised equipment below what was expected to be available. Once that objective was achieved, the interventions were often eased up. However, that meant that the virus was still smouldering. Consequently, it could—and often did—leap into flame again. When that happened, interventions were re-imposed or tightened, with the result that the pattern of cases over time had multiple peaks. This pattern of interventions over time, in which initially a serious effort was made to suppress the virus, which was then slackened when it was successful, and then re-imposed when the virus flared up again, was called 'the hammer and the dance' by an early commentator on COVID-19, Tomas Pueyo.[9]

For example, Figure 5.9 shows the number of daily cases in France from March 2020 to December 2021.

9. Tomas Pueyo, 'Coronavirus: The Hammer and the Dance', 19 March 2020, https://tomaspueyo.medium.com/coronavirus-the-hammer-and-the-dance-be9337092b56.

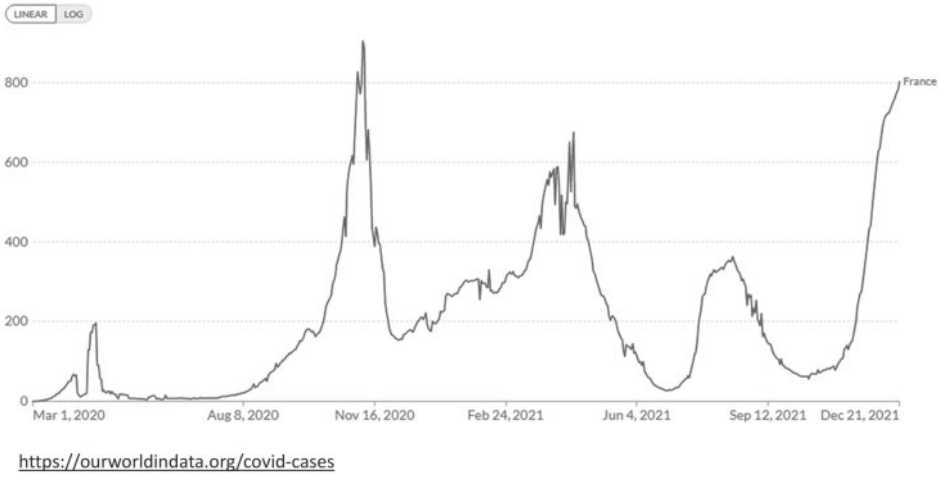

Figure 5.9: Daily new COVID-19 cases in France per million people (seven-day rolling average) to December 2021. Courtesy of Our World in Data.

The peaks represent when the virus flared up. In this graph of actual cases over time, a trough does not represent when the virus dissipated due to herd immunity—as it did in the more theoretical Figures 5.3–5.6—but when government interventions that were tightened or imposed repeatedly in response to the peaks started to be loosened again. Also playing a role was the programme of vaccinations that began at the end of December 2020. In December 2021, Figure 5.9 shows, cases rose sharply because of a new Omicron variant.

Notice that although measures to reduce the rate of growth of cases, hospitalisations, and deaths are often referred to as strategies for 'flattening the curve', they do not, in fact, flatten the curve. The phrase 'flattening the curve' really refers to the attempt to reduce the heights of the peaks so that they do not exceed a nation's capacities, for example for hospital beds. Perhaps instead of the phrase 'flattening the curve', a more accurate phrase should be borrowed from the electric power industry: 'peak shaving'.

How Long Should Someone Who May Have the Disease Be Quarantined?

Some people are impatient with the length of time for which they need to quarantine, especially after returning from a foreign country. For example, in Hong Kong, many travellers returning from overseas were required to quarantine for 21 days at one stage, and then to be tested twice even after quarantine. If the period of time during which someone infected with the disease is infectious averages only several days, why does the quarantine period need to be so long?

The answer has to do with the mathematical concept known as 'fat tails'. Although the average infectious period may be only a few days, there will be variations. Some people will be infectious for longer times, some for shorter times. There will be a distribution of infectious periods, from only three or four days to weeks. Such distributions typically have 'fat tails'—that is, there will be very few people who will have much longer infectious periods than others (i.e., they will be way out on the right-hand tail of the infectious period distribution).

Suppose only one in 10,000 returnees from overseas is infectious for as long as 21 days. That means that if 10,000 people return, there is a good chance one will still be infectious in 21 days, and there is no way to know which one. That is too big a chance to take when a single carrier can ignite a flare-up that can spark exponential spread. If the goal is to ensure no spread, it is prudent to quarantine them all for 21 days. Even supposing that only 1,000 people return, then a chance of one in 10,000 is a chance of one-tenth in 1,000 or still a one-in-ten chance that of those 1,000 one will be infectious in 21 days and ignite exponential spread.

Estimating the Input Parameters to a COVID-19 Prediction Model

Modelling the course of a disease using a predictive simulation requires inputting to the model many assumptions, such as *R0* numbers, hospitalisation rates, death rates, generation time, etc. These are called the *parameters* of the model. They can also include additional numbers, like the percentage of potentially susceptible people who have been vaccinated. And they can include assumptions about what percentage reduction in rates of personal contact will occur when certain mitigating measures are introduced, like school closures.

Because the spread of a disease is exponential in its early stages, decisions need to be made very quickly on whether to adopt policies to clamp down on the rate of growth. These decisions are made with the aid of the mathematical model's projections. For the models to make reasonably accurate predictions, they need reasonably accurate parameters to be input into them. Estimating those parameters in the beginning stages of a new and previously unknown disease, however, is difficult because little data is available.

To help understand the difficulty, consider this dilemma that arose around the end of 2021 and the beginning of 2022. The COVID-19 Omicron variant had just begun to spread rapidly, out-competing the previous variant, Delta, and accounting for the vast majority of COVID-19 cases. Figure 5.10 extends the Figure 5.9 graph of cases in France per million people through the end of the year 2021 and into the beginning of 2022.

Notice how the daily cases of Omicron had, within a space of only two to three weeks, shot up to be much greater than the highest rate before Omicron appeared. The

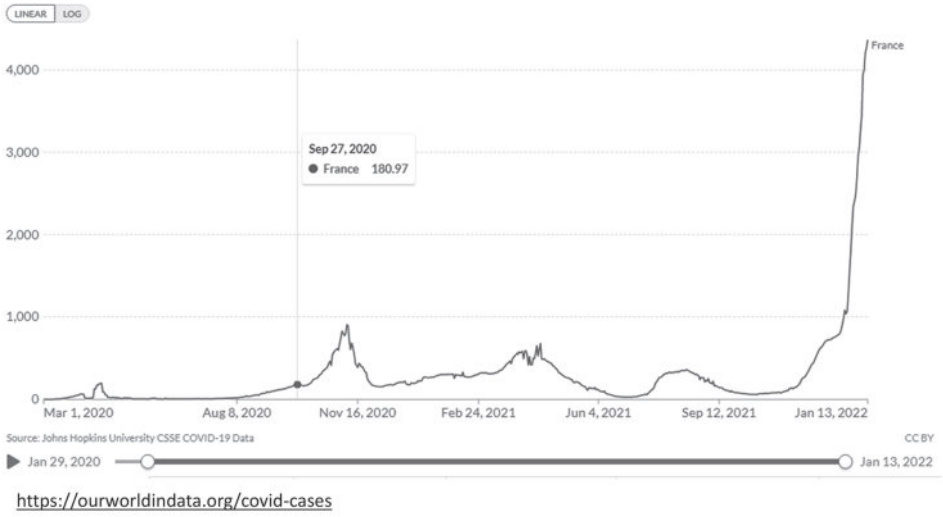

Figure 5.10: Daily new COVID-19 cases in France per million (seven-day rolling average) to mid-January 2022. Courtesy of Our World in Data.

R0 for Omicron was estimated to be as high as 10,[10] on par with the formerly rapidly spreading childhood diseases measles, mumps, and chicken pox (until almost all children were vaccinated for them). An *R0* of 10 implies a doubling of cases approximately every two days.

Early data indicated that the effects of Omicron were milder than previous variants, and it was less likely to require hospitalisation. But because it spread so much more rapidly it might require more hospital beds than previous variants, even though the ratio of hospitalisations to cases was lower. It should be noted that vaccination does not prevent infection, but it lowers the risk of the infected person becoming very sick. With Omicron, the rate of hospitalisation was much lower for those who had been vaccinated. The disease manifested itself as less severe in an infected person who had been vaccinated.

In the United States, there was a desire to estimate what percentage of vaccinated and unvaccinated people who caught Omicron would need hospitalisation. To gather the data needed for this estimate, before Omicron had already spread very widely, was extremely difficult. Hospitalisations lag case discoveries by about two weeks, so the empirical rate of hospitalisation would not be known until at least two weeks after Omicron's onset. Furthermore, Omicron first took hold in regions in the United States where the vaccination rate was high, such as New York, Massachusetts, and New Jersey,

10. Talha Khan Burki, 'Omicron Variant and Booster COVID-19 Vaccines', *The Lancet Respiratory Medicine* 10, no. 2 (2022): e17.

while the spread to regions with lower rates of vaccination, such as rural and Southern regions, took a week or two longer. Hence, it was difficult to estimate the hospitalisation rate for the unvaccinated until as much as a month after Omicron's onset, by which time it might have already peaked.

Estimating a Rate of Growth from Early Data

As an example of how a parameter can be inferred from a small amount of early data, suppose that a researcher, Molly, has only two weeks of data for the daily number of cases of a disease. Let us suppose that the number of cases can be assumed to grow exponentially for at least the next six weeks. How can Molly infer the rate of exponential growth, so that she can extrapolate that rate of growth from the first two weeks to the following six weeks?

Figure 5.11 shows the data for the first 14 days, while Figure 5.12 extends this graph to several possible hypothetical future paths with 14 per cent, 17 per cent, 20 per cent, and 23 per cent daily rates of growth of cases.

Which of these possible growth rates best fits the data we have, which is only for the first 14 days? Figure 5.13 shows 14 per cent, 17 per cent, 20 per cent, and 23 per cent daily rates of growth of cases for the first 14 days.

Figure 5.13 shows the 20 per cent growth rate fits the data best. (A statistical best-fit test would confirm this visual impression.) Therefore, our best-guess projection of future cases in the next six weeks is the 20 per cent growth case in Figure 5.12.

In practice, a modeller will show not only the best-guess projection but an error band with a range of possible projections. Obviously, which growth rate results from the fitting of the 14-day data to a growth rate makes a very big difference. As more data is gathered, the estimate will be revised.

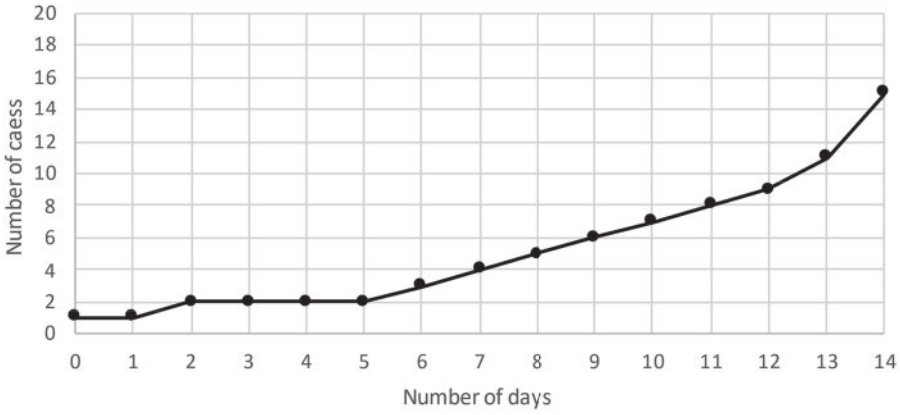

Figure 5.11: Cases for the first 14 days

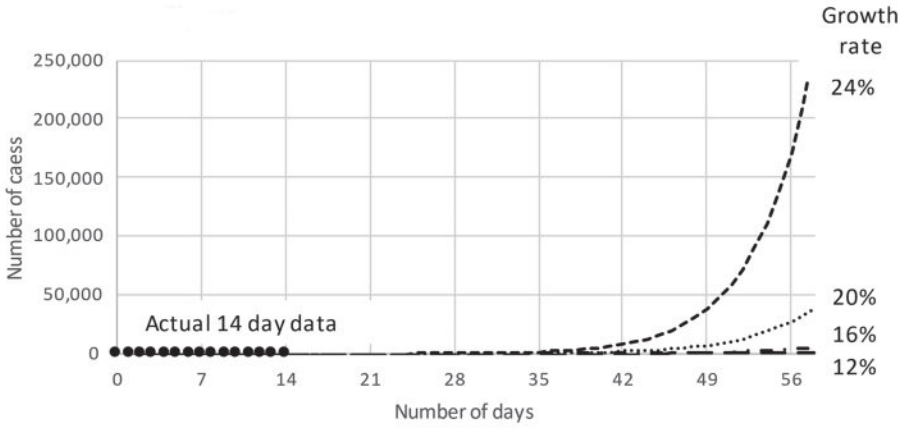

Figure 5.12: Projected cases for 8 weeks

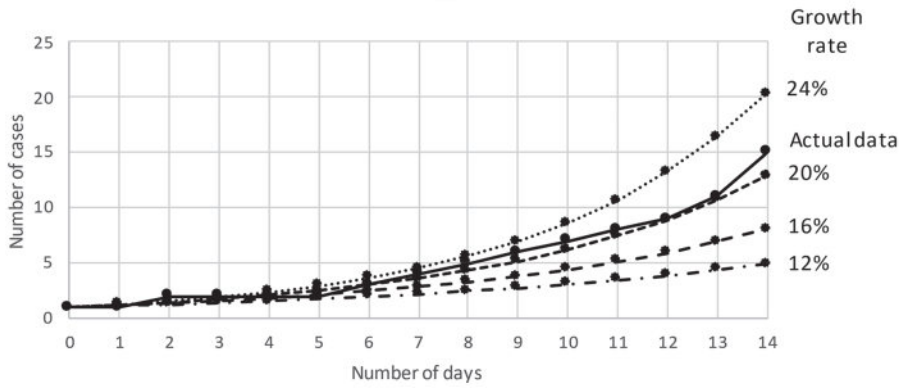

Figure 5.13: Cases for the first 14 days

If the model is an SIR model (which is more appropriate for making predictions of the spread of disease over a period of several months than an exponential model), then a similar approach can be taken to estimating *R0* using early data. Several different *R0*s can be tried, and a simulation run for the first few weeks using each of those hypothetical *R0*s. Whichever of those simulations produces the closest match to the available data, the *R0* that it uses can be adopted for further simulations extended into the future.[11]

Other parameters can be estimated in a similar manner, but a great deal of care is necessary because of the mismatches of data and timing. For example, one cannot

11. The actual methodology used is a little more complicated, but the principle is the same.

estimate the hospitalisation rate or the case fatality rate (the ratio of deaths to symptomatic cases) by dividing the daily or weekly hospitalisations or deaths by the daily or weekly cases for those same days or weeks. Hospitalisations lag the onset of cases—and therefore when the cases are reported—by roughly two weeks and deaths lag by roughly three weeks. Since during the interval between a reported case and a death the number of cases could have ballooned, the death rate could be wildly underestimated.

The Problem of Lag Time in Reporting

In the previous example, in order to project the number of cases into the future, it would be better to know the *R* number, the reproduction rate, on each day. This number can vary from day to day, depending on what social distancing measures are being imposed and adhered to. But the actual *R* number can only be measured after cases become symptomatic and are detected and reported, which can be more than a week after a carrier is infected and becomes infectious. Hence, *R* numbers can only be estimated several days after their impact on disease spread.

To estimate the *R* number in real time, the COVID-19 research team at the University of Hong Kong—which had made some of the first estimates of *R0*, in early 2020—used a novel method, a method that applies particularly well in Hong Kong. Most travellers in Hong Kong use public transportation—buses and the train system known as the MTR—and most of those pay by using an Octopus card, which is swiped on entry to a train station or bus. The Octopus card is also used for small purchases, such as at 7-11 stores.

The level of use of Octopus cards over time is a measure of the level of social mixing. The more the Octopus card is used, the more people are using public transportation and entering stores, and therefore the more they are making contact with one another. The University of Hong Kong researchers calculated the correlation of Octopus card use with past *R* numbers that had been observed after the fact in the population and found that the correlation was strong. Therefore, they estimated real-time *R* numbers by applying that correlation to the level of Octopus card use on a given day and even at a given hour. Using this estimated *R* number, they were able to make instantaneously updated projections of the subsequent spread of disease.

The Hong Kong team was challenged by the rapid spread of Omicron in February and March 2022 to estimate whether and how the spread could be contained. Because of the low percentage of vaccinations among the elderly in Hong Kong, as well as the close quarters in which people live, the challenge was great, especially in residential care homes for the elderly. Their modelling showed that the conclusion was inescapable that the spread would not be fully containable (that the *R* number could not be made to go below 1.5) even with the most stringent control measures that would be practicable in Hong Kong. But because of the rapid spread, more than half the population would be infected and infections would peak by April 2022; however, the risk of resurgences would linger.

Are the Statistics Really What They Seem?

Compounding the difficulty of estimating the parameters of COVID-19 is the fact that the statistics gathered from reporting, recording, or observation are often not what they seem to be. For example, when the Omicron variant broke out, it was found immediately that Omicron spread very quickly but it soon emerged that it caused a lower ratio of hospitalisations to infections, especially in people who were vaccinated. However, as there were so many Omicron infections, the number of hospitalisations was still high. *The New Yorker* magazine, however, noted that a lot of the hospitalisations attributed to Omicron, perhaps as many as half to two-thirds of them, were not due to Omicron at all:

> More than a hundred and fifty thousand Americans are currently hospitalised with the coronavirus—a higher number than at any other point in the pandemic. But that figure, too, is not quite what it seems. Many hospitalized *covid* patients have no respiratory symptoms; they were admitted for other reasons—a heart attack, a broken hip, cancer surgery—and happened to test positive for the virus. There are no nationwide estimates of the proportion of hospitalized patients with "incidental *covid*," but in New York State some forty per cent of hospitalized patients with *covid* are thought to have been admitted for other reasons. The Los Angeles County Department of Health Services reported that incidental infections accounted for roughly two-thirds of *covid* admissions at its hospitals.[12]

Even so, while the hospital admission was for another malady, affliction with the Omicron virus could have been a complicating factor, perhaps enough to drive the person admitted over the threshold for being admitted to a hospital.

It may seem a simple matter to determine how many deaths were caused by COVID-19—just add up all the deaths that were reported to be caused by it. But *The Economist* magazine notes that it is much more complicated than that:

> How many people have died because of the Covid-19 pandemic? The answer depends both on the data available, and on how you define 'because'. Many people who die while infected with SARS-CoV-2 are never tested for it, and do not enter the official totals. Conversely, some people whose deaths have been attributed to Covid-19 had other ailments that might have ended their lives on a similar timeframe anyway. And what about people who died of preventable causes during the pandemic, because hospitals full of Covid-19 patients could not treat them? If such cases count, they must be offset by deaths that did not occur but would have in normal times, such as those caused by flu or air pollution.[13]

12. Dhruv Khullar, 'Do the Omicron Numbers Mean What We Think They Mean', *The New Yorker*, 16 January 2022.
13. *The Economist*, 'The Pandemic's True Death Toll', 2 November 2021, https://www.economist.com/graphic-detail/coronavirus-excess-deaths-estimates.

These complicating factors have caused compilers of statistics to resort to other means to estimate the number of deaths due to COVID-19, and sometimes the number of hospitalisations.

One common means to estimate deaths that can be attributed to the virus is to compare the number of total deaths that occurred while the virus was raging with the number of total deaths 'that would have occurred anyway'. That is, it requires estimating a counterfactual: how many deaths would have occurred if there had been no virus?

For example, Figure 5.14 shows, as a wavy solid line, the number of deaths that would have been expected each week during the years 2018 through 2021 in the United States, extrapolated from the pattern of previous years' weekly deaths.[14] The pattern reflects the fact that more deaths occur during the winter months. The vertical bars are the number of deaths that occurred.

Notice that until April 2020, the number of actual deaths agreed fairly closely with the projected number of deaths from extrapolation. But for most weeks from April 2020 on, the number of actual deaths exceeded the number of projected deaths, in some cases by a wide margin. These excess deaths are very likely attributable to COVID-19. For most countries around the world, excess deaths calculated in this manner do not agree, in many cases not even closely, with the number of deaths reported to have been caused by COVID-19.[15] A March 2022, study in *The Lancet* said that while reported COVID-19 deaths worldwide as of the end of the year 2021 totalled 5.94 million, an estimated 18.2 million died worldwide because of the COVID-19 pandemic as measured by excess mortality—more than three times reported deaths (with a 95 per cent uncertainty interval from 17.1 to 19.6 million).[16] The COVID-19 pandemic has truly been a global tragedy.

Decline in Life Expectancy

Another way to measure the impact of deaths caused by the pandemic is to track life expectancy before the pandemic, and for the years 2020–2021 during the pandemic. It is possible, using mortality data on age at death during a particular year, to calculate life expectancy without having to follow a whole cohort of individuals for their entire lifetimes; in fact, it is much more accurate than following a whole cohort until each of their deaths because life expectancy changes over time.

The method is to calculate the percentage of individuals at each age who died during the year. For example, the data may show that 0.1 per cent of individuals aged zero to one died during the year, while 9 per cent of individuals aged 90 died during the

14. National Center for Health Statistics, 'Excess Deaths Associated with COVID-19', accessed 13 September 2022, https://www.cdc.gov/nchs/nvss/vsrr/covid19/excess_deaths.htm.
15. *The Economist*, 'Tracking Covid-19 Excess Deaths Across Countries', 20 October 2021, https://www.economist.com/graphic-detail/coronavirus-excess-deaths-tracker.
16. COVID-19 Excess Mortality Collaborators, 'Estimating Excess Mortality Due to the COVID-19 Pandemic: A Systematic Analysis of COVID-19-Related Mortality, 2020–21', *The Lancet* 399, no. 10334 (2022): 1513–1536.

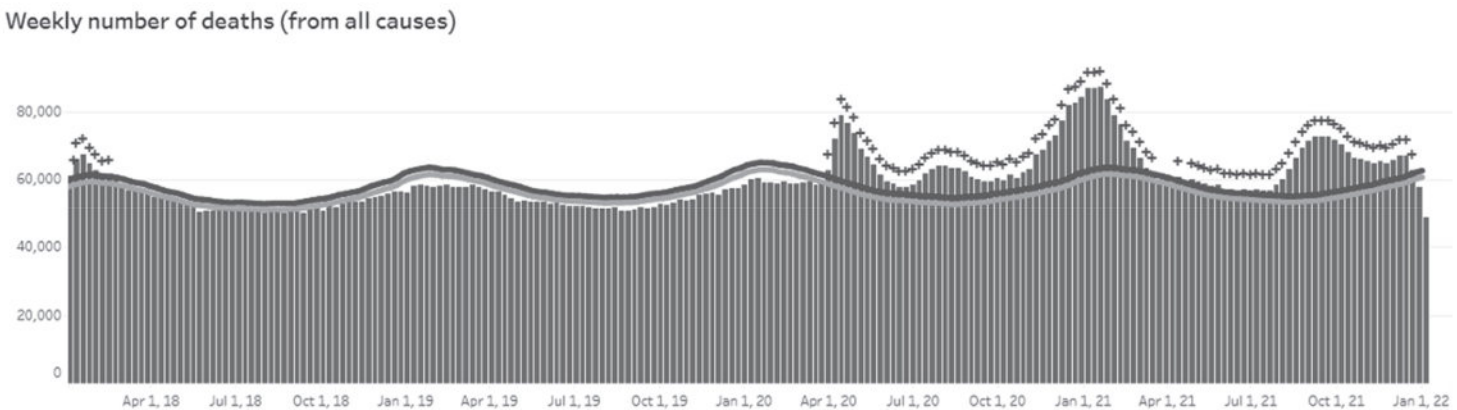

Figure 5.14: Number of actual deaths in the United States versus expected deaths

year. This enables the construction of an entire actuarial table of probabilities of death at each age, from which life expectancy can be calculated.

Not surprisingly, life expectancy declined in several countries. An important study,[17] not yet peer-reviewed at the time of this writing, calculated life expectancy for the three years 2019–2021 (2019 being before the pandemic, and 2020 and 2021 during the pandemic) for the United States and 19 peer countries including Austria, Belgium, Denmark, England and Wales, Finland, France, Germany, Israel, Italy, Netherlands, New Zealand, Northern Ireland, Norway, Portugal, Scotland, South Korea, Spain, Sweden, and Switzerland. The decline in life expectancy over the two years from 2019 to 2021 was by far the largest in the United States. Life expectancy in the United States declined by 2.26 years, from 78.86 years in 2019, before the pandemic, to 76.60 years two years later in 2021, comprising a 1.87-year reduction from 2019 to 2020 and a 0.39-year reduction from 2020 to 2021. By contrast, the other 19 peer countries averaged only a 0.57-year decrease from 2019–2020 and a 0.28-year *increase* from 2020–2021. Among the 19 peer countries, the greatest decline over those two years was 0.93 years, in England and Wales. Life expectancy in the United States was already below that of the peer countries before the pandemic; during the pandemic, the gap increased to more than five years. However, deaths from drug overdoses in the United States also increased by about 30,000 from 2019 to 2021,[18] which cannot be directly attributed to COVID-19; hence, the decline in life expectancy in the United States may slightly overstate the impact of COVID-19.

COVID-19's Last Gasp? Omicron in Shanghai

At the beginning of March 2022, the Omicron variant of COVID-19 began to spread in Shanghai, a city with a population of 25 million in a country of 1.4 billion. China had previously locked down very rapidly once the high transmissibility of the disease and its virulence became clear and henceforth maintained a 'zero-COVID' policy. With the exponential spread of the less virulent Omicron variant in Shanghai, the question arose as to whether the zero-COVID strategy should be maintained (now also called 'dynamic zero' to account for the fact that absolute zero is virtually impossible), or whether something more closely resembling a 'living with COVID' strategy should be initiated. A mathematical model documented by Chinese and US epidemiologists in the journal *Nature* helped to make the decision.[19]

The simple SIR model described earlier in this chapter had assumed that once a person had contracted the disease, they were no longer susceptible to it, at least not

17. Ryan K. Master, Laudan Y. Aron, and Steven H. Woolf, 'Changes in Life Expectancy between 2019 and 2021: United States and 19 Peer Countries', *medRxiv* 1 June 2022, https://doi.org/10.1101/2022.04.05.22273393.
18. National Center for Health Statistics, 'Provisional Drug Overdose Death Counts', accessed 13 September 2022, https://www.cdc.gov/nchs/nvss/vsrr/drug-overdose-data.htm.
19. Jun Cai, Xiaowei Deng, Juan Yang, Kaiyuan Sun, Hengcong Liu, Zhiyuan Chen, et al., 'Modeling Transmission of SARS-CoV-2 Omicron in China', *Nature Medicine* 28 (2022): 1468–1475.

for a long time. But experience showed that people could contract the disease more than once, even within relatively short periods of time. Consequently, the authors of the *Nature* article used an altered version of the model, instead of susceptible-infected-recovered they used a susceptible-latent-infectious-removed-susceptible model to indicate that a person could go from infected and infectious through recovery to susceptible again. The 'latent' phase indicates that an infected person can be asymptomatic in the early stage of the disease.

Running the model produced concerning results, even though Omicron was less deadly than previous COVID-19 variants. The authors reported that: 'We find that the level of immunity induced by the March 2022 vaccination campaign would be insufficient to prevent an Omicron wave that would result in exceeding critical care capacity with a projected intensive care unit peak demand of 15.6 times the existing capacity and causing approximately 1.55 million deaths.' The authors were, nevertheless, confident that continued access to vaccines and antiviral therapies, and implementation of non-pharmaceutical interventions—i.e., lockdowns, social distancing, isolation, and so on—would suffice to prevent overwhelming the healthcare system. Hence, China continued with its lockdown approach in Shanghai, even though the Omicron variant was less deadly and the lockdown was very painful (see Chapter 9).

Host-Parasite Coevolution and the Disease Endgame

As the disease continues to be endemic in the host population and evolves with COVID-19, the virus and its variants become more or less transmissible, and more or less virulent—where 'virulent' means harmful to the host's health or, simply, lethal. It would help to anticipate and plan for the virus 'endgame' if there were a mathematical model to predict how the relationship between virus transmissibility and virulence will evolve over time. Will the disease evolve to become like the common cold, which is highly transmissible but hardly virulent, because the survival rate is virtually 100 per cent? Or will it be more like rabies or tuberculosis, which continue to coexist with a host population seemingly forever? Unfortunately, although there have been more than 6,000 papers published on this question, no definitive answer is available.

It would seem a simple matter of applying Darwin's theory of natural selection to the virus's evolution. What will make virus survival and proliferation more likely—increased virulence of the disease in the host, or decreased virulence? The relevant theoretical relationships in those 6,000 academic papers are expressed in mathematical formulas, but we will describe the basic idea.

From the point of view of the host population, the objective is to make the reproduction number, *R*, as small as possible. For the virus, there is also a reproduction number *R*. In this case, *R* is the number of additional viruses that each virus can spawn and transmit to other hosts. The difference between the host *R* and the parasite's (that is, the virus's) *R* is that while the objective of the human population is to minimise *R*, the objective of the virus is to maximise *R*; i.e., to spread as quickly as possible. Like

the host *R*, the virus's *R* is the product of how fast the virus spreads from one host to another (its transmissivity), and for how many days it can spread from that host; that is, the number of days for which the host is infectious.

For the purposes of this discussion, let us call the latter the duration. The central questions in most of the academic papers are: how long is the duration, and how is its length related to the transmissibility? The longer the duration, the more the virus would be transmitted to other hosts; its *R* would be higher.

Until the 1980s, the prevailing theory was the 'avirulence hypothesis'.[20] This hypothesis assumed that for a virus to be more transmissible, it would have to be less 'virulent'—that is, less lethal—because if it killed the host, that would shorten the duration of the time during which the virus could be transmitted. Therefore, like the common cold, the virus would become milder over time, even if more transmissible. The reason for this assumption was that the less lethal a virus is, the longer the average time for which the host lives, and therefore the longer the time the virus can spend living in the host and transmitting itself to other hosts.

Empirical studies, however—though difficult to perform because the variables are hard to define and measure—were not able to decisively confirm the avirulence hypothesis. In the early 1980s, Anderson and May and others presented another hypothesis,[21] known as the virulence-transmission trade-off. This hypothesis rests on the observation that if a virus has a higher transmission rate, it is likely to be more abundant in a host. Greater abundance means greater cost to the host, which means a higher mortality rate but also a lower daily rate of recovery if the patient does not die—that is, the patient is sick for longer. The higher mortality rate tends to reduce the virus's *R* because it can only live in the host and transmit itself for a shorter time. But the longer recovery time tends to *increase* the virus's *R* because it can live in the sick patient longer.

The virulence-transmission trade-off hypothesis arrives at an optimal transmissivity for the virus given that transmissivity increases mortality, but also increases recovery time for patients who do not die. It is a nice theory, but unfortunately attempts to verify it empirically have stumbled. The problem is twofold: first, it is difficult to gather data for the variables as defined in the model to verify it empirically; and second, other complicating factors can cause the relationship to be different from the result of the theoretical model. The result is that there is no reliable method as yet to predict with any certainty how the virus will evolve over time, and what its transmission rate and virulence will be. There are still hopes for the virulence-transmission trade-off hypothesis, but it has yet to be confirmed at a high level of confidence.

20. S. Alizon, A. Hurford, N. Mideo, and M. Van Baalen, 'Virulence Evolution and the Trade-Off Hypothesis: History, Current State of Affairs and the Future', *Journal of Evolutionary Biology* 22, no. 2 (2009): 245–259; Clayton E. Cressler, David V. McLeod, Carly Rozins, Josée van den Hoogen, and Troy Day, 'The Adaptive Evolution of Virulence: A Review of Theoretical Predictions and Empirical Tests', *Parasitology* 143, no. 7 (2016): 915–930; Miguel A. Acevedo, Forrest P. Dillemuth, Andrew J. Flick, Matthew J. Faldyn, and Bret D. Elderd, 'Virulence-Driven Trade-Offs in Disease Transmission: A Meta-analysis', *Evolution* 73, no. 4 (2019): 636–647.

21. Roy M. Anderson and Robert M. May, 'Coevolution of Hosts and Parasites', *Parasitology* 85, no. 2 (1982): 411–426.

Another phenomenon, however, leads to the conclusion that the virus will weaken over time. As many people in the population contract one variant of the virus or another, and as many people get vaccinated for the virus, their immunity to it and to future variants tends to increase.[22] This has been called 'hybrid immunity'—that is, immunity acquired from both prior infection and vaccination.[23] Therefore, there is a reasonable expectation that COVID-19, while it will be with us for a long time, will gradually become less and less of a problem over time.

Conclusion

The most salient—and dangerous—mathematical feature of the spread of disease is exponential growth. Exponential growth is especially problematic when the disease's transmissivity—its rate of spread, its *R* number—is large. If the rate of spread is not too great, and diligent test-and-trace methods are applied to find and isolate infectious disease carriers before they can spread the disease, the disease can be contained. But if the *R* number is so large that it overwhelms the capacity to test and trace, it can then spread exponentially, catastrophically exceeding hospital and medical care capacities. This was the case with the Omicron variant of COVID-19—though it was at least, fortunately, less deadly than previous variants.

There may be little that can be done to contain a disease that is both highly transmissive and highly virulent, except to ride it out until herd immunity. This is essentially what happened with the black plagues of the Middle Ages, though in those cases the problem was a lack of the medical knowledge needed to contain it. Nevertheless, for many countries of the world, insufficient preparation allowed COVID-19 to spread more than it needed to. With adequate preparation, it could have been contained through better test-and-trace methods and more isolation of the infectious. In advance of a likely future onset of another disease, mathematical modelling of varying hypothetical levels of transmissivity and virulence should be undertaken to determine what levels and types of preparation should be put in place to contain all but the most transmissive, and thus uncontainable, of them. Cost-benefit analyses can be undertaken to determine what levels of defence against the spread of future diseases should be put in place to contain all but—to borrow a measure used to determine how much defence should be put in place against a flood or tsunami—the thousand-year disease. This anticipatory and preventive use of mathematical modelling may be even more essential than using it after a disease strikes.

22. William Hanage, 'After Omicron, This Pandemic Will Be Different', *The New York Times*, 19 January 2022, https://www.nytimes.com/2022/01/19/opinion/omicron-covid-surge.html.
23. Ivan Hung Fan-ngai, 'Beyond Hong Kong's COVID-19 Fifth Wave: Coping with the Coronavirus', *Asia Global Online*, 15 July 2022, https://www.asiaglobalonline.hku.hk/beyond-hong-kongs-covid-19-fifth-wave-coping-coronavirus?utm_source=Asia+Global+Institute&utm_campaign=783e7a6dbd-EMAIL_CAMPAIGN_2020_05_14_04_19_COPY_01&utm_medium=email&utm_term=0_c139173191-783e7a6dbd-381203146.